

IBM Tivoli Storage Productivity Center
Version 4.2.2.3

User's Guide



IBM Tivoli Storage Productivity Center
Version 4.2.2.3

User's Guide



Note:

Before using this information and the product it supports, read the information in "Notices" on page 1219.

This edition applies to version 4, release 2, modification 2, fix pack 3 of IBM Tivoli Storage Productivity Center (product numbers 5608-WB1, 5608-WC0, 5608-WC3, 5608-WC4, 5725-A01, 5725-E18, 5608-WB2, and 5608-WB3) and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces SC27-2338-07.

© Copyright IBM Corporation 2001, 2012.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures ix

Tables xi

Preface xiii

Who should read this guide xiii
Publications xiii
 Accessing publications online xiii
 IBM Redbooks xv
 Translation xv
 Ordering publications xv
 Providing feedback about publications xvi
Conventions used in this guide xvi
Contacting the IBM Support Center xvii
 Reporting a problem xvii

**New for Tivoli Storage Productivity
Center and Tivoli Storage Productivity
Center for Replication Version 4.2.2.3 xix**

**New for Tivoli Storage Productivity
Center and Tivoli Storage Productivity
Center for Replication Version 4.2.2.2 xxi**

**New for Tivoli Storage Productivity
Center and Tivoli Storage
Productivity Center for Replication
Version 4.2.2.1. xxiii**

**New for Tivoli Storage Productivity
Center and Tivoli Storage Productivity
Center for Replication Version 4.2.2. xxv**

Chapter 1. Getting started 1

Chapter 2. Product overview 11

Introducing IBM Tivoli Storage Productivity Center 11
 Architecture 12
 IBM Tivoli Storage Productivity Center summary 13
 Licenses for Tivoli Storage Productivity Center 14
 Data Manager 22
 Disk Manager 24
 Tape Manager 26
 Fabric Manager 26
 Graphical user interface 27
Tivoli Storage Productivity Center for Replication 49
 Tivoli Storage Productivity Center for Replication
 Two Site Business Continuity 49
 Tivoli Storage Productivity Center for Replication
 Three Site Business Continuity 51

Tivoli Storage Productivity Center for Replication
integration with Tivoli System Automation
 Application Manager 53
Tivoli Integrated Portal 53
Key concepts 54
 Starting and stopping the Tivoli Storage
 Productivity Center GUI 54
 Single sign-on 58
 Launch and launch in context 60
 User roles 71
 User tasks 74
 Data collection 77
 Data replication 79
 Alerts 109
 Provisioning planner 111
 Reporting. 115
 Fabric and zones 115
 Agents 119
 Security 126
 SAN Volume Controller topics. 128
Scenarios 136
 Monitoring and reporting on I/O utilization 137
 Storage provisioning using workflows 140
 Archive and delete files 144
 Create a profile with a filter 147
 Creating a performance threshold within an
 alert 150
 Determine application performance problems 153
 Monitor space on a remote server 157
 Identifying performance hot spots 161
 Retiring underutilized storage systems 163
 Tivoli Storage Productivity Center for
 Replication scenarios 165

Chapter 3. Administering 181

Administering DB2 181
 Backing up the Tivoli Storage Productivity
 Center database 181
 Starting the Control Center 182
 Using the command line on UNIX and Linux 182
 Manually starting DB2 on Windows 182
 Manually stopping DB2 on Windows 183
 Monitoring DB2 183
Adding a Tivoli Storage Productivity Center user
group name to Tivoli Storage Productivity Center
for Replication 184
Accessing the Tivoli Storage Productivity Center
for Replication GUI 184
Backing up and recovering all the Tivoli Storage
Productivity Center for Replication 185
Administering Tivoli Storage Productivity Center 186
 Checking the Tivoli Storage Productivity Center
 for Replication version 186
 Starting and stopping the Tivoli Storage
 Productivity Center GUI. 186

Stopping the Tivoli Storage Productivity Center GUI	190
Viewing and managing services	190
Starting and stopping the Tivoli Storage Productivity Center services	198
Starting and stopping the Tivoli Integrated Portal server.	202
Checking Tivoli Storage Productivity Center status	202
Increasing memory allocation	203
Setting timeout values for the Device server	205
Changing passwords	206
Checking for a fully qualified host name	211
Granting local administrative privileges to a domain account	213
Importing SSL certificates for VMware	213
SQL access to Tivoli Storage Productivity Center views	214
Repository Copy tool.	218
Service tool: collecting information	220
Administering Tivoli Integrated Portal	225
Changing the user authentication method	225
Changing the LTPA token expiration for single sign-on	242
Administering data sources.	242
Checking the health of your data sources	242
Manually changing the Windows service logon for the agent	243
CIM agents	243
Data agents and Storage Resource agents	250
In-band Fabric agents	274
Out-of-band Fabric agents	276
Tivoli Storage Productivity Center servers	278
Storage systems	284
VMware data source	292
Managing schedules, jobs, and runs	295
Job Management - Recommendations	302

Chapter 4. Managing storage resources 305

Configure Devices.	305
Configure storage subsystems	306
Configure fabrics and switches	314
Configure computers	317
Configure tape libraries	319
Discovering storage resources	321
Discovery overview	321
Discovering computers in the Windows domain, NAS, and SAN File System.	329
Discovering computers	330
Discover storage subsystem, tape library, and fabric information with CIM agents	330
Discovering fabrics with Storage Resource agents	332
Discovering storage subsystems and switches	333
Collecting data	335
Collecting storage statistics (probes)	335
Collecting storage statistics from Tivoli Storage Productivity Center servers (TPC server probes)	340
Collecting storage usage and trending information (scans)	343

Determining the availability of storage resources (pings)	364
Data collection groups	368
Working with groups of resources	371
Collecting performance data	375
Monitoring switch performance	375
Monitoring storage-subsystem performance	376
Monitoring tape libraries	381
Viewing performance data	381
Working with alerts	383
Creating alerts	383
Configuring alert destinations	394
Deleting alerts	395
Disabling an alert	395
Editing alerts	396
Renaming alerts	397
Viewing alerts	398
Viewing alert logs	401
Viewing tape library alerts	402
Checking storage configurations	402
Using configuration history.	402
Using configuration analysis	407
Planning and modifying storage configurations	417
SAN Planner overview	417
Using the SAN Planner	422
Working with storage resource group sessions	434
Optimizing storage configurations	434
Storage Optimizer overview	435
Using the Storage Optimizer	436
Enforcing storage policies	444
Setting storage usage quotas	446
Storage usage for Network Appliance quotas	466
Setting file constraints	471
Using filesystem extension	480
Scheduling script-based actions	493
Archive and backup functions.	498
Working with storage resource groups	503
Storage resource group management	503
Create or Edit a Storage Resource Group	504
Working with scripts	507
Running scripts with alerts.	507
Running scripts using scheduled actions	508
How scripts are run	510
How the server and agent resolve script names	512
Using the topology viewer	513
The topology viewer	513
Arranging entities in a group	527
Checking the status of entities and groups.	527
Creating custom groups	536
Defining properties for grouping entities	537
Displaying data paths from one endpoint in the topology to another	537
Displaying membership in zones and zone sets	546
Displaying detailed information about entities	547
Displaying storage resource connections	547
Filtering entities based on health status.	548
Labeling entities	548
Starting element managers	548
Launching SAN Planner.	550
Starting Storage Optimizer	550
Locating storage resources in the topology	550

Managing Storage Subsystems	551	Starting and stopping IBM Tivoli Storage Productivity Center for Replication	635
Managing Zone Control	551	Verifying that components are running	637
Pinning entities	551	Starting the IBM Tivoli Storage Productivity Center for Replication GUI	638
Printing information from the topology viewer	552	Identifying the version of IBM Tivoli Storage Productivity Center for Replication	639
Refreshing views	552	Backing up and restoring IBM Tivoli Storage Productivity Center for Replication configuration data	639
Removing entities from the database	552	Managing management servers	643
Removing deleted entities	553	Management servers	643
Selecting entities in graphical and table views	553	Ports	644
Showing Alerts	553	SNMP alerts	646
Showing the monitoring status of entities	554	Setting up a standby management server	649
Working with tape libraries in the topology viewer	554	Reinstalling the primary server during an active session	650
Managing disk storage	555	Reconnecting the active and standby management servers	651
Working with workload profiles	556	Performing a takeover on the standby management server	651
Working with provisioning profiles	558	Configuring SNMP	651
Working with thin provisioning profiles	563	Adding SNMP managers	652
Managing storage subsystems	565	Changing the standby management server port number	652
Managing virtual disks	573	Changing the client port number	652
Managing volumes	582	Managing storage systems	653
Managing fabrics	588	Ports	653
Working with zones	588	Storage systems	655
Working with zone sets	591	Storage connections	656
Working with zone aliases	594	Protected volumes	659
Managing Tape Libraries	596	Site awareness	660
Discovering tape libraries	597	Adding a storage connection	660
Viewing tape library information	597	Removing a storage connection	662
Monitoring tape libraries	598	Removing a storage system	662
Managing DS8000 storage systems using element managers	598	Modifying the location of storage systems	663
Configuring the DS8000 element manager	599	Modifying storage connection properties	663
Configuring other element managers for DS8000	604	Refreshing the storage system configuration	664
Managing other storage systems using element managers	605	Setting volume protection	664
Adding element managers using CIMOM discovery	606	Restoring data from a journal volume	665
Modifying other element managers	608	Managing host systems	665
Starting other element managers	609	Adding a host system connection	666
Testing the connection to other element managers	609	Modifying a host system connection	666
Adding user-defined element managers for other storage systems	610	Removing a host system connection	667
Removing user-defined element managers for other storage systems	610	Removing a session from a host system connection	667
Launch and launch in context	610	Managing logical paths	668
Inbound launch in context feature	611	Viewing logical paths	668
Outbound launch and launch in context features	618	Adding logical paths	668
External tools	622	Adding logical paths using a CSV file	669
Adding external tools	622	Removing logical paths	670
Modifying external tools	623	Setting up data replication	670
Starting external tools	624	Sessions	670
Removing external tools	624	Data replication	676
Working with invoices	625	Session commands	707
Planning for chargeback	626	Site awareness	722
Setting up invoices	627	Preserve Mirror option	723
Generating an invoice	628	Creating sessions and adding copy sets	725
Viewing and printing invoices	628	Using the Metro Mirror heartbeat	731
Sample Invoices	629	Exporting copy set data	734
		Importing copy set data	734
Chapter 5. Managing replication	635		
Administering	635		

Modifying the location of session sites	735
Removing sessions	735
Removing copy sets	735
Migrating an existing configuration to Tivoli Storage Productivity Center for Replication	736
Practicing disaster recovery.	739
Practice volumes	739
Practicing disaster recovery for a Metro Mirror Failover/Failback with Practice session.	740
Practicing disaster recovery for a Global Mirror either Direction with two-site Practice session	740
Practicing disaster recovery for a Global Mirror Failover/Failback with Practice session.	741
Practicing disaster recovery for a Metro Global Mirror Failover/Failback with Practice session	741
Monitoring health and status	742
Viewing the health summary	742
Viewing SNMP alerts.	742
Viewing sessions	743
Viewing storage system details	761
Viewing storage connection details	762
Viewing volume details	762
Viewing logical paths.	762
Viewing console messages	763
Security	763
Users and groups	763
User roles	764
Adding the IBM Tivoli Storage Productivity Center for Replication Administrator role to the IBM Tivoli Storage Productivity Center Superuser group	766
Granting access privileges for a user.	767
Viewing access privileges for a user	767
Modifying access privileges for a user	767
Removing access privileges for a user	768
Chapter 6. Reporting	769
Choosing a reporting type	769
Choosing a report category.	771
What can I find out using reports?	777
Report considerations and limitations	781
Working with reports.	784
Generating and viewing reports	784
Sorting report columns	785
Printing reports or charts	785
Including charts in reports	786
Drilling down a report row.	789
Navigating tabs	789
Trending with historical data	790
User-defined online reports.	793
Generating reports	793
Using reporting groups	801
Data Manager reports	804
Data Manager for Databases reports.	933
Data Manager for Chargeback reports	961
Disk Manager reports	961
Fabric Manager reports	973
Tape Manager reports	980
System reports	980
Generating and viewing system reports	980
System reports - Data	981

System reports - Data Manager for Databases	993
System Reports - Disk Manager	999
System reports - Fabric Manager	1016
Batch reports	1021
Creating batch reports	1021
Editing batch reports	1032
Viewing job information for a batch report	1033
Displaying log files for a batch report	1033
Displaying batch reports from the saved directory	1033
Viewing CSV batch reports with spreadsheet applications	1034
Exporting and importing settings that you define	1034
User ID (saved) reports.	1034
Creating a user-defined report	1035
Displaying user-defined reports	1035
Editing user-defined reports	1035
Rollup reports.	1035
Generating rollup reports	1036
Rollup Reports - Asset	1037
Rollup Reports - Database Asset.	1066
Rollup Reports - TPC Server Storage Space	1072
Rollup Reports - Database Capacity	1088
Data source reports - overview	1093
Data source reports - CIMOM agents	1093
Storage tier reports - overview	1097
Running storage tier reports	1097
Tivoli Storage Productivity Center Monitoring Agent reports	1130
Appendix A. Reference	1137
Alerts and schedules associated with monitoring groups and templates	1137
Supported Fabric Agent Types	1143
Supported storage systems providing full disk encryption and solid-state drives	1145
Frequently Asked Questions	1145
General and Migration Information.	1145
Data Manager	1147
Fabric Manager	1147
Disk Manager	1150
Arguments and window locations available with the inbound launch in context feature	1151
Launch in context parameters	1151
Launch in context task parameters	1153
Available windows	1154
Alerts	1156
Triggering conditions for alerts	1158
Alert suppression actions	1171
Triggered actions for alerts	1172
Alert display colors	1173
Performance data.	1175
General performance guidelines	1175
Performance thresholds.	1176
Performance metrics.	1186
Protocols and standards	1208
Web Based Enterprise Management	1208
Storage Management Initiative Specification	1209
Service Location Protocol	1210
Simple Network Management Protocol	1210

IBM Tivoli Storage Productivity Center		Notices	1219
Universal Agent	1211	Trademarks	1220
Fibre Channel Methodologies of Interconnects	1212		
Appendix B. Accessibility features for Tivoli Storage Productivity Center	1213	Glossary	1223
Appendix C. Accessibility features for Tivoli Storage Productivity Center for Replication	1215	Index	1229

Figures

1. Tivoli Storage Productivity Center components	13	34. Snapshot selection panel	404
2. Tivoli Storage Productivity Center Main Window	28	35. Icons and colors of the change overlay	406
3. User roles in Tivoli Storage Productivity Center	71	36. The Create Analyzer window	409
4. A SAN Volume Controller node	129	37. Topology overview	515
5. Example of a SAN Volume Controller in a fabric	130	38. Minimap	516
6. Physical disks and managed disks (MDisks)	131	39. Computer group	517
7. MDisk group	133	40. Data Path Explorer view	539
8. Managed disk groups and VDIs	135	41. Connections	547
9. High-level architecture	141	42. Highlighted entity	550
10. Selecting Federated repositories	228	43. Pinned entity	551
11. Manage repositories	228	44. Storage-Subsystem Details page	568
12. Adding a new repository	229	45. Create Virtual Disk wizard: "Define the virtual disks" page	575
13. Enter values for a new repository	230	46. Virtual Disks page	578
14. Add base entry to realm	230	47. Virtual-Disk Details page	580
15. Enter values for a new repository continued	231	48. Volumes page	582
16. Configuring Federated repository	232	49. Sample Invoice — OS Storage Usage By User	630
17. Configuring Federated repository continued	232	50. Sample Invoice — OS Disk Capacity by Computer	631
18. Selecting local operating system	239	51. Sample Invoice — Storage Usage by Database User	632
19. Configuring local operating system	240	52. Sample Invoice — Total Size by Database-Tablespace	633
20. Selecting local operating system for available realm definition	241	53. Chart Generation	787
21. Master and subordinate server architecture	279	54. Tabs	790
22. Editing the CIMOM discovery job (Options tab)	331	55. Trending with Historical Data	791
23. Editing the CIMOM discovery job (When to Run tab)	332	56. History Chart for File System Used Space	793
24. Checking storage subsystems	332	57. Selection Resources window	798
25. File Filters tab, pop-up menu	355	58. Usage Reporting—Access Load—Access Time—Chart Network Wide	893
26. Create Condition window	356	59. Usage Reporting—Access Load—Modification Time—Chart Network Wide	893
27. Create Group Window	356	60. Usage Reporting—Files—File Size Distribution—File Count Chart Tab	897
28. File Filter window	357	61. Managed SAN overview report	1131
29. Alert History page	398	62. Managed SAN Health report	1132
30. Detail for Alert page	399	63. TPC Server overview report	1133
31. Enabling or disabling the alert overlay	399	64. TPC Server overview report	1134
32. Using the pop-up menu to view alerts in the table view	400		
33. Viewing alerts by using the Alert tab in the table view	401		

Tables

1.	Locations of publications for Tivoli Storage Productivity Center and related products . . . xiv	42.	Metro Mirror with Practice commands 712
2.	Tivoli Storage Productivity Center summary of functions 14	43.	Global Mirror commands 713
3.	Available functions in each Tivoli Storage Productivity Center license 15	44.	Global Mirror with Practice commands 715
4.	Available reports in each IBM Tivoli Storage Productivity Center license 19	45.	Metro Global Mirror commands 716
5.	License upgrade scenarios for IBM Tivoli Storage Productivity Center 21	46.	Metro Global Mirror with Practice commands 719
6.	Functions that require components to be licensed and installed 72	47.	Session status icons 743
7.	GUI nodes that require special roles or licenses 72	48.	Volume role symbols 743
8.	Session type summary 79	49.	Data copying symbols 744
9.	Roles and authorization levels 127	50.	Session states 746
10.	Managed disk group states 133	51.	Detailed status messages for Participating and Non-Participating role pairs. 759
11.	Capacities of the cluster given extent size 134	52.	Predefined monitoring groups, templates, schedules, and alerts for storage subsystems 1138
12.	Virtual disk states 136	53.	Predefined monitoring groups, templates, schedules, and alerts for fabrics 1140
13.	Role-pair states at the completion of each step (participating and nonparticipating role pairs) 175	54.	Predefined monitoring groups, templates, schedules, and alerts for computers 1141
14.	Role-pair states at the completion of each step (participating and nonparticipating role pairs) 177	55.	Predefined monitoring groups, templates, and schedules for tape libraries ¹ 1143
15.	Role-pair states at the completion of each step (participating and nonparticipating role pairs) 179	56.	Supported agent types for Switch and Fabric Functions 1143
16.	Discovery types, their data sources, and top-level entities that are discovered 322	57.	Interop Namespaces for CIM Agents for switches and storage systems 1146
17.	Information collected by agent 325	58.	Supported agent types for Switch and Fabric Functions 1147
18.	Supported SMI agent profiles and subprofiles 328	59.	Available windows for inbound launch in context 1154
19.	Icons and colors of the change overlay 405	60.	Triggering conditions and related error messages for computer alerts 1158
20.	Policies used by the configuration analysis feature 409	61.	Triggering conditions and related error messages for file system alerts 1159
21.	Maximum volume name lengths for storage systems 426	62.	Triggering conditions and related error messages for directory alerts 1160
22.	Descriptions of policy types and their triggering conditions 444	63.	Triggering conditions and related error messages for instance alerts 1160
23.	LUN Provisioning Results 485	64.	Triggering conditions and related error messages for database-tablespace alerts 1161
24.	LUN Provisioning Results 491	65.	Triggering conditions and related error messages for table alerts 1162
25.	Group locations in the topology viewer 518	66.	Triggering conditions and related error messages for storage-subsystem alerts 1163
26.	Topology status indicators and their icons 527	67.	Triggering conditions and related error messages for fabric alerts 1167
27.	Calculation of health status for a storage entity 532	68.	Triggering conditions and related error messages for switch alerts 1168
28.	Categories of operational status for storage entities. 533	69.	Triggering conditions and related error messages for endpoint device alerts. 1169
29.	Storage system default ports 645	70.	Triggering conditions and related error messages for hypervisor alerts 1169
30.	Session state change traps 647	71.	Triggering conditions and related error messages for NAS alerts 1169
31.	Configuration change traps 647	72.	Alert colors 1173
32.	Suspending-event notification traps 648	73.	Array thresholds 1177
33.	Communication-failure traps 648	74.	Controller thresholds. 1181
34.	Management Servers state-change traps 648	75.	Port thresholds. 1182
35.	Storage system default ports 654		
36.	Session type summary 676		
37.	Basic HyperSwap commands 708		
38.	FlashCopy commands. 708		
39.	Snapshot session commands 709		
40.	Snapshot group commands 709		
41.	Metro Mirror commands. 710		

Preface

IBM® Tivoli® Storage Productivity Center is a storage infrastructure management software product that can centralize, automate, and simplify the management of complex and heterogeneous storage environments.

Who should read this guide

The audience of this publication is administrators or users who have already installed the product and are now using IBM Tivoli Storage Productivity Center. This publication provides an overview of how to use the product and detailed information about various topics.

Administrators and users must be familiar, or become familiar, with the following concepts:

- SAN
- Tivoli Storage Productivity Center
- IBM DB2
- Simple Network Management Protocol (SNMP)
- IBM Tivoli Enterprise Console®

Publications

This section lists publications in the IBM Tivoli Storage Productivity Center library and other related publications. It also describes how to access publications online, how to order publications, and how to submit comments on publications.

The publications are available from the IBM Publications Center at <http://www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss>.

Accessing publications online

Publications for this product and other related products are available online.

Information about installing, configuring, upgrading, and uninstalling IBM Tivoli Storage Productivity Center and related products is available online. Use Table 1 on page xiv to view and download these publications. Translated documents are available for some products.

Table 1. Locations of publications for Tivoli Storage Productivity Center and related products

Product	Online location (see Note)
IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication	<p>http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp. In the left navigation pane, click IBM Tivoli Storage Productivity Center.</p> <p>To obtain PDF documents, click IBM Tivoli Storage Productivity Center > Printable documentation.</p> <p>To view previous versions of the IBM Tivoli Storage Productivity Center Information Center, in the left navigation pane click Previous Versions and click the version that you want to view (for example, 4.1.1, 4.1, or 3.4.1).</p>
IBM Tivoli Storage Productivity Center for Replication for System z®	<p>http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp. In the left navigation pane, click IBM Tivoli Storage Productivity Center for Replication for System z.</p> <p>To obtain PDF documents, click IBM Tivoli Storage Productivity Center for Replication for System z > Printable documentation.</p> <p>To view previous versions of the IBM Tivoli Storage Productivity Center for Replication Information Center, in the left navigation pane click Previous Versions and click the version that you want to view (for example, 4.1.1, 4.1, or 3.4.1).</p>
IBM WebSphere® Application Server	<p>http://publib.boulder.ibm.com/infocenter/wasinfo/v6r1/index.jsp</p>
IBM System Storage® Productivity Center	<p>http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp. In the left navigation pane, click System Storage Productivity Center.</p> <p>To obtain PDF documents, click System Storage Productivity Center > Printable documentation.</p> <p>To view previous versions of the IBM System Storage Productivity Center Information Center, in the left navigation pane, click Previous Versions and click the version that you want to view (for example, 1.3.2, 1.3.1, or 1.3).</p>
IBM System Storage DS3000, IBM System Storage DS4000®, or IBM System Storage DS5000	<p>http://www.ibm.com/support/entry/portal/</p>
IBM System Storage DS6000™	<p>http://publib.boulder.ibm.com/infocenter/dsichelp/ds6000ic/index.jsp</p>
IBM System Storage DS8000®	<p>http://publib.boulder.ibm.com/infocenter/dsichelp/ds8000ic/index.jsp</p>
IBM System Storage DS® Open Application Programming Interface publications	<p>http://www.ibm.com/support/entry/portal/Troubleshooting/Hardware/System_Storage/Storage_software/Other_software_products/CIM_Agent_for_DS_Open_(API)/</p> <p>Use these publications for information about how to install, configure, and use the CIM agent.</p>

Table 1. Locations of publications for Tivoli Storage Productivity Center and related products (continued)

Product	Online location (see Note)
IBM System Storage SAN Volume Controller	http://publib.boulder.ibm.com/infocenter/svc/ic/index.jsp
IBM Storwize® V7000	http://publib.boulder.ibm.com/infocenter/storwize/ic/index.jsp
IBM Storwize V7000 Unified	http://publib.boulder.ibm.com/infocenter/storwize/unified_ic/index.jsp
IBM XIV® Storage System	http://publib.boulder.ibm.com/infocenter/ibmxiv/r2/index.jsp
IBM DB2® Database for Linux, UNIX, and Windows	http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp
IBM System Storage N series	http://www.ibm.com/systems/storage/network/redbooks.html For more information about IBM System Storage N series, see http://www.ibm.com/systems/storage/network/software/ .
Note: If you print PDF documents on other than letter-sized paper, in the Print window select Fit to Printable Area in the Page Scaling field. This option is available when you click File > Print . Fit to Printable Area ensures that the full dimensions of a letter-sized page print on the paper that you are using.	

IBM Redbooks

The IBM Redbooks® are books on specialized topics.

You can order publications through your IBM representative or the IBM branch office serving your locality. You can also search for and order books of interest to you by visiting the IBM Redbooks home page at <http://www.redbooks.ibm.com/redbooks>.

Translation

Translated publications are available from the information center which is available in certain translated languages. It is displayed in the language that is appropriate for the browser locale setting.

When a locale does not have a translated version, the information center is displayed in English, which is the default language. Translations of the PDFs are available when the information center is translated.

Contact your service representative for more information about the translated publications and whether translations are available in your language.

Ordering publications

Information is provided for the ordering of IBM publications on the Internet or by telephone.

You can order many IBM publications online at <http://www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss>.

You can also order by telephone. In the United States and Canada, call 800-879-2755. In other countries, contact your IBM service representative.

Providing feedback about publications

Your feedback is important to help IBM provide the highest quality information. You can provide comments or suggestions about the documentation from the IBM Tivoli Storage Productivity Center Information Center.

Go to the information center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp> and click the **Feedback** link at the bottom of the information center Welcome page or topic pages.

Conventions used in this guide

This section provides information about the conventions used in this publication.

This publication uses several conventions for special terms and actions, and operating system-dependent commands and paths.

The following typeface conventions are used in this publication:

Bold

- Lower-case and mixed-case commands that appear with text
- Command options that appear with text
- Flags that appear with text
- Graphical user interface (GUI) elements (except for titles of windows and dialogs)
- Names of keys

Italic

- Variables
- Values you must provide
- New terms
- Words and phrases that are emphasized
- Titles of documents

monospace

- Commands and command options in examples
- Flags that appear on a separate line
- Code examples and output
- Message text
- Names of files and directories
- Text strings you must type, when they appear within text
- Names of Java methods and classes
- HTML and XML tags also appear like this, in monospace type

For syntax notation, these conventions are used:

- # is the prompt for the root user on UNIX platforms.
- Uppercase and lowercase characters do matter for UNIX and Linux. Type in commands exactly as shown.

Contacting the IBM Support Center

You can find support information for IBM products from various sources.

For support for IBM Tivoli Storage Productivity Center, you can contact the IBM Support Center in one of the following ways:

- Go to the IBM Tivoli Storage Productivity Center technical support website at http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center_Standard_Edition.

To receive future support notifications, sign in under **Notifications**. You are required to enter your IBM ID and password. After you are authenticated, you can configure your subscription for Tivoli Storage Productivity Center technical support website updates.

- Customers in the United States can call 1-800-IBM-SERV (1-800-426-7378).
- For international customers, go to the Tivoli Storage Productivity Center technical support website for customer support telephone numbers.

You can also review the *IBM Software Support Handbook*, which is available on our website at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html>.

The support website offers extensive information, including a guide to support services; frequently asked questions (FAQs); and documentation for all IBM Software products, including Redbooks and white papers. Translated documents are also available for some products.

When you contact the IBM Support Center, be prepared to provide identification information for your company so that support personnel can readily assist you. Company identification information might also be needed to access various online services available on the website. See “Reporting a problem.”

Reporting a problem

This topic provides a list of what information you should have ready when you encounter a problem.

Have the following information ready when you report a problem:

- The IBM Tivoli Storage Productivity Center version, release, modification, and service level number.
- The communication protocol (for example, TCP/IP), version, and release number that you are using.
- The activity that you were doing when the problem occurred, listing the steps that you followed before the problem occurred.
- The exact text of any error messages.

New for Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication Version 4.2.2.3

There are no new features in IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication version 4.2.2.3.

For limitations and known issues for 4.2.2.3, go to the Tivoli Storage Productivity Center support site at http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center_Standard_Edition. In the **Search support and downloads** field, search for the document *Limitations and known issues for Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication 4.2.2.3*.

New for Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication Version 4.2.2.2

Use this information to learn about new features and enhancements in Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication Version 4.2.2.2. This topic highlights the changes since Tivoli Storage Productivity Center Version 4.2.2.1.

Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication 4.2.2.2 add the following new features, functions, and enhancements.

Tivoli Storage Productivity Center

GPT partition type

You can use Tivoli Storage Productivity Center to monitor disks that use GUID partition table (GPT) partition style for New Technology File Systems (NTFS) in addition to the master boot record (MBR) partition style.

Cisco Proxy SMI Agent Support

You can use Tivoli Storage Productivity Center with the Cisco Data Center Network Manager (DCNM) SMI-S Server as a CIMOM agent.

LVM2 support on RHEL5

You can use Tivoli Storage Productivity Center to monitor file systems that are mounted on volume groups and logical volumes that are managed by a Logical Volume Manager (LVM).

VMware 5 Hypervisor and Tivoli Storage Productivity Center SRA-only support

You can use Tivoli Storage Productivity Center to monitor VMware 5 Hypervisor servers. The Storage Resource agent can be installed on vSphere 5 virtual machines.

IBM Tivoli Monitoring Agent for Tivoli Storage Productivity Center support for Linux on System z.

You can use the Tivoli Storage Productivity Center Monitoring agent for Linux on System z.

Tivoli Storage Productivity Center for Replication

Additional support for space-efficient volumes in Global Mirror with Practice sessions

You can use extent space-efficient volumes as copy set volumes for Global Mirror with Practice sessions for System Storage DS8000 6.3 or later.

For more information, see “Copy sets” on page 671.

New features

The following features are new for Tivoli Storage Productivity Center for Replication for System z Version 4.2.2.2.

Reflash After Recover option for Global Mirror Failover/Failback with Practice sessions

You can use the **Reflash After Recover** option with System Storage DS8000 Version 4.2 or later. Use this option to create a FlashCopy® replication between the I2 and J2 volumes after the recovery of a

Global Mirror Failover/Failback with Practice session. If you do not use this option, a FlashCopy replication is created only between the I2 and H2 volumes.

For more information, see “Session properties” on page 748.

No Copy option for Global Mirror with Practice and Metro Global Mirror with Practice sessions

You can use the **No Copy** option with System Storage DS8000 Version 4.2 or later. Use this option if you do not want the hardware to write the background copy until the source track is written to. Data is not copied to the I2 volume until the blocks or tracks of the H2 volume are modified.

For more information, see “Session properties” on page 748.

StartGC H1->H2 command for Global Mirror sessions

You can use the StartGC H1->H2 command with TotalStorage Enterprise Storage Server® Model 800, System Storage DS8000, and System Storage DS6000. This command establishes Global Copy relationships between site 1 and site 2, and begins asynchronous data replication from H1 to H2.

For more information, see “Session commands” on page 707.

Language pack

Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication Version 4.2.2.2 provide a Language Pack that you can use to see messages, online help, and text in a language other than English. The Language Pack is available for download with the fix pack.

New for Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication Version 4.2.2.1

Use this information to learn about new features and enhancements in Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication Version 4.2.2.1. This topic highlights the changes since Tivoli Storage Productivity Center 4.2.2.

Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication 4.2.2.1 add the following new features, functions, and enhancements:

IBM Storwize V7000 Unified

Storwize V7000 Unified is a virtualized storage system that includes Storwize V7000 and Storwize V7000 File Module. Storwize V7000 Unified is designed to consolidate block and file workloads into a single storage system for simplicity of management, reduced cost, highly scalable capacity, performance, and high availability. Storwize V7000 Unified also offers improved efficiency and flexibility through built-in solid state drive (SSD) optimization, thin provisioning and non-disruptive migration of data from existing storage. The system can virtualize and reuse existing disk systems offering a greater potential return on investment.

Data for Storwize V7000 Unified storage systems is collected, monitored, displayed, and reported in Tivoli Storage Productivity Center and data replication is supported by Tivoli Storage Productivity Center for Replication.

System Storage DS8000 R6.2

Starting with System Storage DS8000 R6.2, you are no longer required to access the DS Storage Manager through the Tivoli Storage Productivity Center GUI. The DS Storage Manager can be accessed from any location that has network access using a web browser. For a list of browsers that are supported for DS8000, see the DS8000 information center.

New for Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication Version 4.2.2

Use this information to learn about new features and enhancements in Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication Version 4.2.2. This topic highlights the changes since Tivoli Storage Productivity Center 4.2.1.

Tivoli Storage Productivity Center

Tivoli Storage Productivity Center 4.2.2 adds the following new features, functions, and enhancements:

IBM Tivoli Storage Productivity Center Select

The Select Edition includes the same functions and features available in IBM Tivoli Storage Productivity Center Standard Edition:

- Disk, fabric, tape, and data management functions
- Storage system performance monitoring, fabric performance monitoring, chargeback, and database monitoring
- Analytical functions, such as:
 - configuration analysis
 - storage area network (SAN) Planner
 - Storage Optimizer
 - configuration history

This product supports the IBM, XIV Storage System, IBM Storwize V7000, IBM System Storage DS3000, IBM System Storage DS4000, and IBM System Storage DS5000 as stand-alone devices, or when attached to the IBM System Storage SAN Volume Controller. This product also supports any storage devices that are attached to Storwize V7000.

IBM Tivoli Storage Productivity Center for Disk Select

Tivoli Storage Productivity Center for Disk Midrange Edition has been renamed to Tivoli Storage Productivity Center for Disk Select. This product provides the same function of the Tivoli Storage Productivity Center for Disk Midrange Edition:

- Basic disk, fabric, tape, and data management functions
- Storage system performance monitoring

This product supports the XIV system, Storwize V7000, DS3000, DS4000, and DS5000 as stand-alone devices, or when attached to the SAN Volume Controller. This product also supports any storage devices that are attached to Storwize V7000.

Storage tier reports

Storage tier reports help you decide where to place accessed data for the best performance from your system, based on whether the data is accessed frequently or infrequently. Infrequently accessed data can be placed on lower-cost storage systems and frequently accessed data can be placed on faster performing storage systems.

To create the reports, you must have Tivoli Storage Productivity Center Standard Edition or Tivoli Storage Productivity Center Select installed. You

must also have Tivoli Common Reporting installed. These reports are supported for IBM System Storage SAN Volume Controller and Storwize V7000. To generate the reports, you must set the back-end subsystem pool properties, and then run a performance monitor job for the SAN Volume Controller or Storwize V7000. You can then display the reports using Tivoli Common Reporting.

Tivoli Storage Productivity Center reports

For SAN Volume Controller clusters and Storwize V7000 systems, you can drill down from storage subsystems and expose managed disks, managed disk groups, and virtual disks. For Storwize V7000 systems, you can also drill down from the internal disks node to view the internal storage on the system or generate a Storage Subsystem Disks report to view the internal storage on the system.

SAN Planner

The SAN Planner now supports workload profiles and replication planning for XIV systems.

Tivoli Storage Productivity Center for Replication

Tivoli Storage Productivity Center for Replication 4.2.2 supports IBM XIV Storage System. You can use the following session types for the XIV system:

Snapshot

Snapshot is a new session type that creates a point-in-time copy (snapshot) of a volume or set of volumes without having to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.

Metro Mirror Failover/Failback

Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 kilometers apart. You can use failover and failback to switch the direction of the data flow.

Global Mirror Failover/Failback

Global Mirror is a method of asynchronous, remote data replication that operates between two sites that are over 300 kilometers apart. You can use failover and failback to switch the direction of the data flow.

Tivoli Storage Productivity Center for Replication support for the XIV system includes the following new features.

Support for volume nickname

For the XIV system sessions, you can provide the volume ID or the volume nickname as a parameter value when you add or remove copy sets by using the command line interface (CLI) commands **mkcpset** and **rmcpset**.

In addition, you can include the XIV system volume ID or the volume nickname in a comma-separated value (CSV) file that is used to import copy set information. You can import the CSV file by using the **importcsv** command or the Tivoli Storage Productivity Center for Replication graphical user interface (GUI). CSV files that are exported from Tivoli Storage Productivity Center for Replication for XIV system sessions include the volume nickname rather than the volume ID. CSV files are exported by using the **exportcsv** command.

New CLI commands

The following CLI commands are new. For more information about new

and updated CLI commands for Tivoli Storage Productivity Center for Replication 4.2.2, see the IBM Tivoli Storage Productivity Center Command-line Interface Reference.

cmdsnapgrp

Use the **cmdsnapgrp** command to run a specific action against a snapshot group that is in an XIV system Snapshot session. A snapshot group is a grouping of snapshots of individual volumes in a consistency group at a specific point in time.

lssnapgrp

Use the **lssnapgrp** command to view snapshot groups that are in an XIV system Snapshot session.

lssnapgrpactions

Use the **lssnapgrpactions** command to specify the session and snapshot group name that you want to view available actions for.

lssnapshots

Use the **lssnapshots** command to view snapshots that are in a snapshot group in an XIV system session.

Chapter 1. Getting started

Use this task list as a getting started guide for installing and using IBM Tivoli Storage Productivity Center. The tasks are in sequential order to help you learn how the product works and to get started administering and managing your storage environment.

Complete the following tasks as shown and click the included links to learn more about the capabilities of Tivoli Storage Productivity Center. Follow the tasks as outlined: planning, installing, configuring, administering, managing storage devices, and reporting. Troubleshooting contains procedures for debugging any problems that you might encounter.

Task	Description	When and how often must I perform this task?
Determine prerequisites and plan for installation.	<p>The first installation task is planning for Tivoli Storage Productivity Center. The time you invest in planning directly affects your ease of implementation. Planning includes:</p> <ul style="list-style-type: none">• Ensuring that your systems meet hardware and software requirements• Determining where to deploy Storage Resource agents• Ensuring that ports are open for use by Tivoli Storage Productivity Center components• Granting privileges to the user ID used to install the Tivoli Storage Productivity Center components• Setting up for LDAP Support	Perform this task before you install Tivoli Storage Productivity Center.
Install Tivoli Storage Productivity Center.	<p>The second task is installing Tivoli Storage Productivity Center.</p> <p>You can install all the Tivoli Storage Productivity Center components using typical installation or custom installation. It is recommended that for a production environment you use custom installation.</p> <p>Tivoli Storage Productivity Center has the following installable components:</p> <ul style="list-style-type: none">• Database schema• Data server• Device server• Graphical user interface (GUI)• Command-line interface (CLI)• Storage Resource agent• IBM Tivoli Integrated Portal	Perform this task after you have completed planning and ensuring that you meet the installation prerequisites. You can also install components separately depending on your needs, so you might perform a portion of the installation process a number of times after the initial installation. For example, you can install Storage Resource agents at any time on the computers whose storage you want to monitor.

Task	Description	When and how often must I perform this task?
Start and Configure Tivoli Storage Productivity Center.	<p>Perform this task after installing Tivoli Storage Productivity Center to configure it according to the standards within your organization, including:</p> <ul style="list-style-type: none"> • Starting the Tivoli Storage Productivity Center GUI. • Running the Configuration Utility as a guide to help you through the steps required to configure and manage Tivoli Storage Productivity Center for your environment. • Using the Help menu to access the online help system, launch the product information center, and access information about the appropriate levels of software, firmware, and hardware that is needed to run the product. • Assigning roles to individuals who use the product. • Entering licensing information. • Setting SNMP and email alerts. • Entering retention settings for log files, resource history, and removed resources. • Defining rules for generating email addresses of quota violators. • Scheduling history aggregation and agent upgrades. • Setting up Novell NetWare and NAS filer access. • Setting up IBM TotalStorage Enterprise Storage Server or SAN Volume Controller access. • Using configuration history to analyze historical data that identifies possible problems with a storage area network (SAN) configuration. • Using configuration analysis to determine whether an existing SAN configuration complies with predefined best practices. 	<p>Typically, you perform this step immediately after installing the product. Once Tivoli Storage Productivity Center is configured, you do not have to perform this step again unless changes within your organization require it.</p>

Task	Description	When and how often must I perform this task?
Administer DB2 and Tivoli Storage Productivity Center.	<p>Once Tivoli Storage Productivity Center is installed and configured, perform the following administrative tasks as needed:</p> <ul style="list-style-type: none"> • Authorizing users. • Managing Agent Manager. • Administering data sources (agents): CIMOM, Data, Storage Resource, Inband Fabric, Out of Band Fabric, Storage Subsystems, TPC Servers, VMWare VI. • Starting and stopping Tivoli Storage Productivity Center services • Increasing the memory allocation and checking the status of Tivoli Storage Productivity Center components. • Setting timeout values for the Device server. • Changing passwords for Tivoli Storage Productivity Center components • Configuring the launch of the IBM Tivoli Storage Productivity Center for Replication GUI and other web GUIs or local executable files. • Collecting information with the service tool. 	<p>Perform administrative tasks at any time during the use of Tivoli Storage Productivity Center. These tasks are meant to provide you with the tools to maintain a successful implementation of the product during its lifecycle.</p>
Configure storage devices for monitoring	<p>Before you can use Tivoli Storage Productivity Center to monitor the storage devices in your environment, you must configure them within the user interface. Configuring devices includes the following tasks:</p> <ul style="list-style-type: none"> • Discovering storage devices. • Including storage devices in monitoring groups (optional). • Defining data collection schedules. • Determining the conditions that must occur before you are alerted to a problem related to the devices you are monitoring. <p>There are two methods for configuring storage devices in Tivoli Storage Productivity Center</p> <ul style="list-style-type: none"> • Use the Configure Devices wizard to guide you through the steps of configuring devices. The wizard is available under the IBM Tivoli Storage Productivity Center node in the navigation tree. • Manually configure devices using the corresponding nodes in the navigation tree. <p>The following rows in the table provide more information about these tasks.</p>	<p>Perform this task when you want to configure storage devices for monitoring by Tivoli Storage Productivity Center. You can configure devices that are already known to Tivoli Storage Productivity Center but not included in a probe schedule, or configure devices that have not yet been discovered.</p>

Task	Description	When and how often must I perform this task?
Configure storage devices: running discovery schedules	<p>Perform discovery to detect the storage devices and network topology within your environment. Discovery schedules collect basic information about the storage devices, including computers, CIMOMs, storage subsystems, fabrics, NetWare trees, filers, clusters, and tape libraries.</p> <p>To discover storage devices, use either of the following methods:</p> <ul style="list-style-type: none"> • Use the discovery step in the Configure Devices wizard to automate the detection of storage devices. • Go to the Administrative Services > Discovery node in the navigation tree and manually define a schedule for the device type that you want to detect. 	<p>You must run discovery schedules before you can collect more detailed information about storage resources. Once discovery has run, use data collection schedules such as probes, scans, and pings to collect more detailed information about the discovered storage devices.</p> <p>See Discover resources in your system for more information.</p>
Configure storage devices: defining monitoring groups for the related storage devices in your environment	<p>Organize related storage devices into groups (optional). Use groups when you want to perform the same action against multiple devices. For example, if you create a group that contains all the storage systems for the Sales division within your organization, you can include that group in a data collection schedule. The schedule then collects information about the all the storage systems in the Sales group without requiring you to add them individually to a schedule.</p> <p>As new devices are introduced into the environment, you can add them to existing groups. The same actions that are already associated with the group are then applied to the new devices.</p> <p>To create groups, use either of the following methods:</p> <ul style="list-style-type: none"> • Use the monitoring group step in the Configure Devices wizard to automate the grouping of devices. • Use the Monitoring > Groups facility within each of Tivoli Storage Productivity Center managers to define a group that corresponds to that manager. For example, use Disk Manager to create storage subsystem monitoring groups; use Fabric Manager to create fabric groups; use Data Manager to create computer, file system, directory, user, and OS user monitoring groups. 	<p>Create monitoring groups before you do the following:</p> <ul style="list-style-type: none"> • Define data collection schedules (probes, scans, and pings) to collect information about storage devices. • Define Data Manager quotas and constraints to enforce storage policies. <p>Creating groups is not required, but they are meant to help you better organize your storage devices for monitoring in a large environment.</p> <p>See the following topics for more information:</p> <ol style="list-style-type: none"> 1. "Data collection groups" on page 368 2. "Working with groups of resources" on page 371

Task	Description	When and how often must I perform this task?
<p>Configure storage devices: determining the information you want to gather about the devices in your environment and defining data collection schedules.</p>	<p>Determining what information you want to gather about the storage devices within your enterprise is critical to helping you implement a storage management strategy.</p> <p>Once you have determined the types of information you want to collect about your storage devices, you can use data collection schedules to gather that information. Use the following schedules to collect data about storage devices:</p> <ul style="list-style-type: none"> • Probes: collect detailed statistics about storage devices. • Scans: collect data about the usage and trending of your storage consumption. • Pings: collect information about the availability (uptime) of computers. • Storage Subsystem Performance Monitors: collect statistics about the performance of storage subsystems. • Fabric Performance Monitors: collect statistics about the performance of fabrics. • TPC Server probes: gather the information that has been collected by multiple Tivoli Storage Productivity Center servers. <p>To create data collection schedules, use either of the following methods:</p> <ul style="list-style-type: none"> • For creating probes and pings, you can use the data collection step in the Configure Devices wizard. • Go to the Monitoring > <i>schedule_name</i> node available in each manager, where <i>schedule_name</i> represents the schedule type (for example, probe, scan, ping, performance monitor). For example, to create a scan, go to the Data Manager > Monitoring > Scan node in the navigation tree. 	<p>Define data collection schedules after you determine the type of data you want to collect. You must collect data about storage devices before you can perform the other tasks within Tivoli Storage Productivity Center, such as viewing reports, enforcing storage policies, setting alerts, using the topology viewer, and managing storage subsystems.</p> <p>After you define a data collection schedule, you do not need to redefine that schedule in the future unless you want to add or remove devices. The schedule runs at the times you specified in its definition.</p> <p>See the following topics for more information:</p> <ul style="list-style-type: none"> • Collect data from system resources • “Collecting storage statistics (probes)” on page 335 • “Collecting storage usage and trending information (scans)” on page 343 • “Determining the availability of storage resources (pings)” on page 364 • “Collecting performance data” on page 375 • “Managing fabrics” on page 588 • “Collecting storage statistics from Tivoli Storage Productivity Center servers (TPC server probes)” on page 340

Task	Description	When and how often must I perform this task?
<p>Configure storage devices: determining the conditions that must occur before you are alerted to a problem within your storage environment.</p>	<p>Determining when and how you are alerted to conditions within your storage environment is important to helping you maintain and administer storage devices. Tivoli Storage Productivity Center provides an extensive list of conditions from which you can choose. See “Triggering conditions for alerts” on page 1158 for a list of triggering conditions for storage devices.</p> <p>When a condition is met and alert is triggered, you can choose to be notified through an SNMP trap, IBM Tivoli Enterprise Console event, login notification, event log (Windows), syslog (UNIX), or email.</p> <p>To associate alert conditions with storage devices, use either of the following methods:</p> <ul style="list-style-type: none"> • Use the Configure Devices wizard to guide you through the steps for associating alert conditions with the storage devices that you want to monitor. • Go to the Alerting node available in each manager to manually associate a specific type of alert condition with storage devices. For example, to associate alert conditions with computers, go to the Data Manager > Alerting > Computer Alerts node in the navigation tree. 	<p>Define alerts:</p> <ul style="list-style-type: none"> • After defining data collection schedules • Based on the information gathered by data collection schedules <p>Alerts are checked whenever a probe or scan schedule is collects data about storage devices. Computer, fabric, switch, and endpoint alerts are triggered by probes. Directory and storage subsystem alerts are triggered by scans. File system free space alerts are triggered by both probes and scans. All other file system alerts are triggered only by probes.</p> <p>After you define an alert, you do not need to redefine it in the future (unless you want to change it). Defining an alert is a one-time only task.</p> <p>See “Working with alerts” on page 383 for more information.</p>

Task	Description	When and how often must I perform this task?
<p>Manage storage resources: determining the storage policies that you want to enforce.</p>	<p>Enforcing storage policies within your organization can be critical to ensuring the proper and most cost-effective use of your storage devices. Use the following policy management functions to set policies for devices that are monitored by Tivoli Storage Productivity Center:</p> <ul style="list-style-type: none"> • Use quotas to define limits on the amount of storage that a user or a group of users can consume. • Use Network Appliance Quotas to import, view, and report on the NetApp Quotas defined on NAS filers. • Use constraints to define the acceptable and unacceptable file types, file sizes, and file owners for a computer or a set of computers in your environment. • Use file system extensions to automatically extend file systems when utilization reaches a specified threshold. • Use scheduled actions to run scripts against storage devices on a schedule that you define. 	<p>A quota runs according to the schedule that you specified. Because quotas look at the information gathered by scans to determine quota violations, you must schedule quotas to run after the appropriate scan. For example, if you want to restrict the amount of space consumed by users on specific computers, run a scan on those computers to collect usage information before running the quota to enforce that usage restriction.</p> <p>Constraints are automatically checked whenever a scan is run against the storage devices to which the constraint applies.</p> <p>Schedule file system extensions after running probes of the related storage devices and file systems.</p> <p>Run scheduled actions immediately, once at a specified time, or repeatedly according to a schedule you specify.</p> <p>See the following topics for more information:</p> <ol style="list-style-type: none"> 1. Creating and enforcing storage policies 2. “Setting storage usage quotas” on page 446 3. “Setting file constraints” on page 471 4. “Using filesystem extension” on page 480 5. “Scheduling script-based actions” on page 493
<p>Manage storage devices: working with storage systems and fabrics.</p>	<p>Use the following functions to manage the storage systems and fabrics that are monitored by Tivoli Storage Productivity Center:</p> <ul style="list-style-type: none"> • Use the Storage Optimizer identify performance bottlenecks and recommend changes to improve performance. • Use the SAN Planner to assist in the end-to-end planning of fabrics, hosts, storage controllers, storage pools, volumes, paths, ports, zones, and zone sets. • Define storage resource groups to logically group related storage entities such as hosts, storage subsystems, fabrics, and switches. • Define element managers for storage subsystems. • Manage virtual disks: assign host ports, create virtual disks, delete virtual disks, add a managed disk to a managed-disk group, view virtual-disk information. • Manage volumes: assign host ports, create volumes, delete volumes, view information about volumes. 	<p>You can manage storage subsystems at any time after collecting information about them through discovery and monitoring jobs.</p> <p>See the following topics for more information:</p> <ul style="list-style-type: none"> • “Optimizing storage configurations” on page 434 • “Planning and modifying storage configurations” on page 417 • “Working with workload profiles” on page 556 • “Managing storage subsystems” on page 565 • “Managing virtual disks” on page 573 • “Managing volumes” on page 582

Task	Description	When and how often must I perform this task?
Manage storage devices: charging for storage usage	Use the storage usage information gathered by data collection schedules to generate invoices that charge back for storage usage.	Generate invoices whenever you want to charge users or groups of users in your organization for their storage usage. See “Working with invoices” on page 625 for more information.
Manage storage devices: setting up and managing replication	Use Tivoli Storage Productivity Center for Replication to manage the advanced copy services for storage systems, including setting up management servers, adding storage connections and paths, creating sessions and copy sets, setting up data replication, performing planned maintenance, and recovering from a disaster. Use IBM Tivoli Storage Productivity Center for Replication Two Site Business Continuity to obtain continuous availability and disaster recovery solutions by using point-in-time replication, which includes FlashCopy and Snapshot, and continuous replication, which includes Metro Mirror and Global Mirror. Use IBM Tivoli Storage Productivity Center for Replication Three Site Business Continuity to obtain continuous availability and disaster recovery solutions by using point-in-time replication, which includes FlashCopy and Snapshot, and continuous replication, which includes Metro Mirror, Global Mirror, and Metro Global Mirror to secondary and tertiary sites.	Access the Tivoli Storage Productivity Center for Replication user interface at any time to set up management servers and connections to storage systems. See Chapter 5, “Managing replication,” on page 635 for more information.
View information about storage devices: using the topology viewer	Use the topology viewer to see a graphical representation of the physical and logical storage resources that have been discovered in your environment. You can view the health status, performance, and other details about storage resources through this function. The topology viewer consists of two views, a graphical view and a table view, organized vertically with the graphical view on top of the table view. The table view shows the same information as the graphical view, but in a tabular format.	You can use the topology viewer at any time after running monitoring jobs to collect the information about your storage environment. See “Using the topology viewer” on page 513 for more information.

Task	Description	When and how often must I perform this task?
View information about storage devices: using any of the following report types: predefined, user-defined, batch, and rollup reports.	<p>Use reports to view the information that is gathered by data collection and quota schedules.</p> <p>Consider the following steps when determining how to best use Tivoli Storage Productivity Center reports:</p> <ol style="list-style-type: none"> 1. Determine how you want to generate and view reports by selecting a reporting type that best meets your needs. 2. Identify which reports contain the most important information about your storage environment. 3. Run the appropriate data collection schedules. 4. Generate the reports you want to view. 	<p>Perform this task after the appropriate data collection schedules have run and you want to view information about your storage devices.</p> <p>See the following topics for more information:</p> <ol style="list-style-type: none"> 1. "Choosing a reporting type" on page 769 2. "Choosing a report category" on page 771 3. "What can I find out using reports?" on page 777 4. "Working with reports" on page 784
Troubleshooting	<p>Use troubleshooting tips to resolve any problems you might have with Tivoli Storage Productivity Center.</p>	<p>Perform troubleshooting tasks when you encounter problems while working with the product. Troubleshooting information is meant to help you with resolving issues that might occur during the use of the product.</p>

Chapter 2. Product overview

IBM Tivoli Storage Productivity Center provides a set of policy-driven automated tools for managing storage capacity, availability, events, performance and assets in your enterprise environment, including NetWare, NAS, SSA, TotalStorage Enterprise Storage Server, and Microsoft Cluster Server technologies, as well as RDBMSs such as Oracle, Sybase SQL Server, Microsoft SQL Server, and DB2 UDB. Tivoli Storage Productivity Center provides storage management from the host and application to the target storage device. It provides disk and tape subsystem configuration and management, performance management, SAN fabric management and configuration, and usage reporting and monitoring.

Tivoli Storage Productivity Center can help you identify, evaluate, control and predict your enterprise storage management assets. Because it is policy-based, it can detect potential problems and automatically make adjustments based on the policies and actions that you define. For example, it can notify you when your system is running out of disk space or warn you of impending storage hardware failure. By alerting you to these and other issues related to your stored data, it enables you to prevent unnecessary system and application downtime.

Tivoli Storage Productivity Center:

- Simplifies the management of storage infrastructures
- Manages, configures, and provisions SAN-attached storage
- Monitors and tracks performance of SAN-attached devices
- Monitors, manages, and controls (through zones) SAN fabric components
- Manages the capacity utilization and availability of file systems and databases

For select IBM hardware orders only, IBM Tivoli Storage Productivity Center Limited Edition, limited to 2 TB storage managed, is provided at no charge, and with one year software maintenance. IBM Tivoli Storage Productivity Center Standard Edition, includes support for TS3500 and limited support for 3494 tape libraries.

Introducing IBM Tivoli Storage Productivity Center

This section provides an overview of IBM Tivoli Storage Productivity Center.

IBM Tivoli Storage Productivity Center provides a set of policy-driven automated tools for managing storage capacity, availability, events, performance and assets in your enterprise environment, including NetWare, NAS, SSA, TotalStorage Enterprise Storage Server, and Microsoft Cluster Server technologies, as well as RDBMSs such as Oracle, Sybase SQL Server, Microsoft SQL Server, and DB2 UDB. Tivoli Storage Productivity Center provides storage management from the host and application to the target storage device. It provides disk and tape subsystem configuration and management, performance management, SAN fabric management and configuration, and usage reporting and monitoring.

Tivoli Storage Productivity Center can help you identify, evaluate, control and predict your enterprise storage management assets. Because it is policy-based, it can detect potential problems and automatically make adjustments based on the policies and actions that you define. For example, it can notify you when your

system is running out of disk space or warn you of impending storage hardware failure. By alerting you to these and other issues related to your stored data, it enables you to prevent unnecessary system and application downtime.

Tivoli Storage Productivity Center:

- Simplifies the management of storage infrastructures
- Manages, configures, and provisions SAN-attached storage
- Monitors and tracks performance of SAN-attached devices
- Monitors, manages, and controls (through zones) SAN fabric components
- Manages the capacity utilization and availability of file systems and databases

Architecture

The IBM Tivoli Storage Productivity Center consists of several key components. This topic identifies these components and shows how they are related.

Data server

This component is the control point for product scheduling functions, configuration, event information, reporting, and graphical user interface (GUI) support. It coordinates communication with and data collection from agents that scan file systems and databases to gather storage demographics and populate the database with results. Automated actions can be defined to perform file system extension, data deletion, and Tivoli Storage Productivity Center backup or archiving, or event reporting when defined thresholds are encountered. The Data server is the primary contact point for GUI user interface functions. It also includes functions that schedule data collection and discovery for the Device server.

Device server

This component discovers, gathers information from, analyzes performance of, and controls storage subsystems and SAN fabrics. It coordinates communication with and data collection from agents that scan SAN fabrics and storage devices.

Database

A single database instance serves as the repository for all Tivoli Storage Productivity Center components.

Agents

Storage Resource agent, CIM agents, and Out of Band fabric agents gather host, application, storage system, and SAN fabric information and send that information to the Data Server or Device server.

Tip: Data agents and Fabric agents are supported in V4.2. However, no new functions were added to those agents for that release. For optimal results when using Tivoli Storage Productivity Center, migrate the Data agents and Fabric agents to Storage Resource agents.

GUI Use the graphical user interface (GUI) to send and receive information for all Tivoli Storage Productivity Center components.

CLI Use the command-line interface (CLI) to issue commands for key Tivoli Storage Productivity Center functions.

Tivoli Integrated Portal

The Tivoli Storage Productivity Center installation program includes IBM Tivoli Integrated Portal. Tivoli Integrated Portal is a standards-based architecture for Web administration. Tivoli Integrated Portal enables developers to build administrative interfaces for IBM and independent

software products as individual plug-ins to a common console network. The installation of Tivoli Integrated Portal is required to enable single sign-on for Tivoli Storage Productivity Center. *Single sign-on* is an authentication process that you can use to enter one user ID and password to access multiple applications. Single sign-on integrates with launch in context features so you can move smoothly from one application to another application.

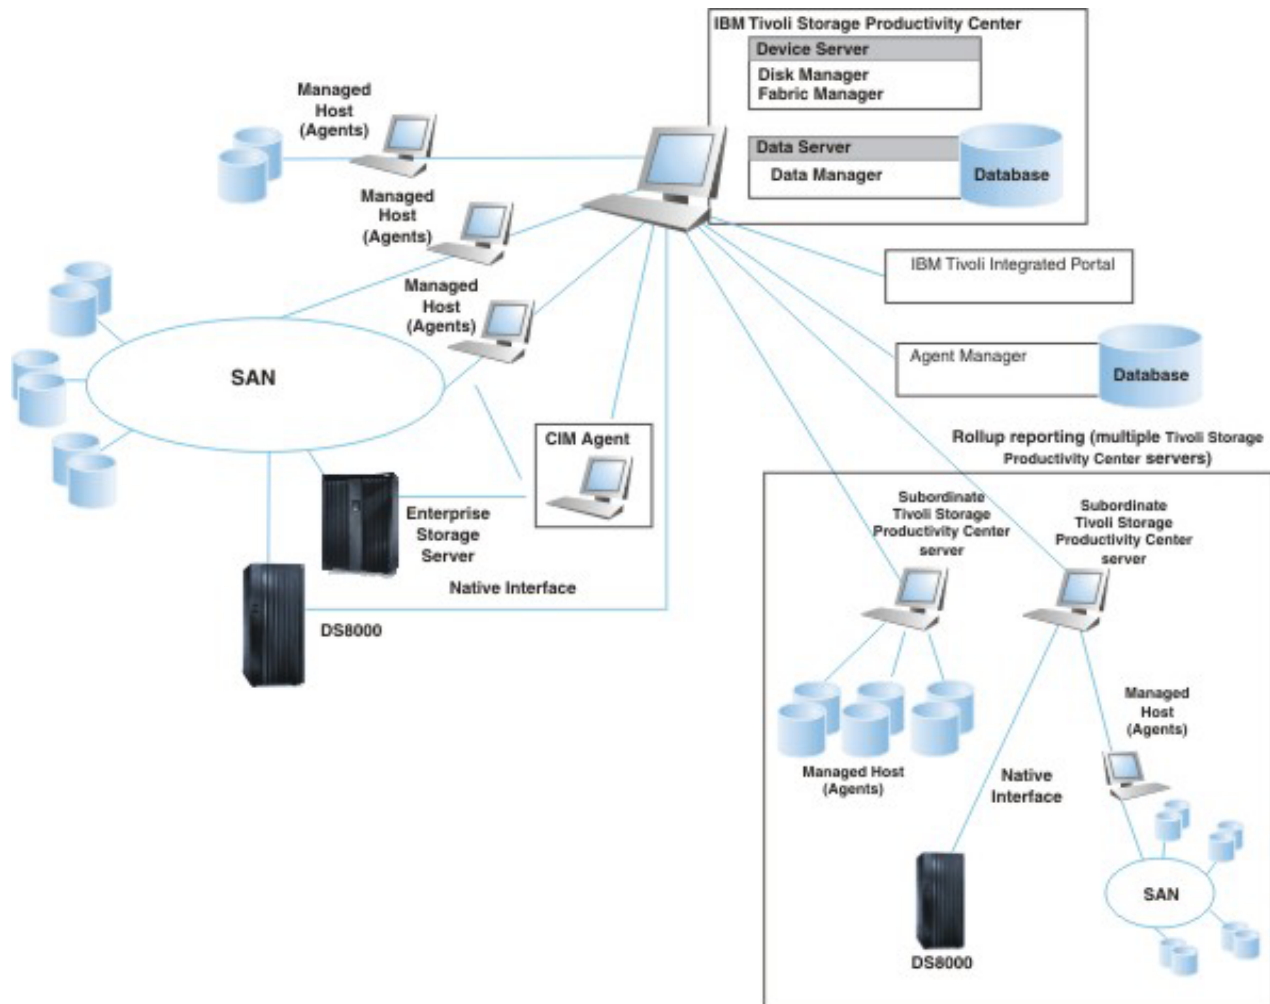


Figure 1. Tivoli Storage Productivity Center components

IBM Tivoli Storage Productivity Center summary

This topic provides a summary of the functionality provided by Tivoli Storage Productivity Center.

The Tivoli Storage Productivity Center functions are shown in Table 2 on page 14.

Table 2. Tivoli Storage Productivity Center summary of functions

Data management	Disk management	Fabric management	Tape management
Host-centric: <ul style="list-style-type: none"> • Discovery of storage resources • Monitoring of storage resources • Filesystem extension • Enterprise-wide reporting Application-centric: <ul style="list-style-type: none"> • Monitoring of relational databases DB2, Oracle, SQL Server, Sybase • Chargeback 	For storage subsystems: <ul style="list-style-type: none"> • Discovery of storage resources • Monitoring of storage resources • Configuration (for example, creating volumes) • Performance management 	For fabrics: <ul style="list-style-type: none"> • Discovery of storage resources • Monitoring of storage resources • Configuration (for example, zoning) • Performance management 	For tape libraries: <ul style="list-style-type: none"> • Discovery of storage resources • Monitoring of storage resources

Licenses for Tivoli Storage Productivity Center

The following product licenses are available for IBM Tivoli Storage Productivity Center: Basic, Data, Disk, Select, and Standard Editions. Each license determines the functions that are accessible in the user interface. This topic lists the functions that are included with each license to help you determine which one best meets your storage management needs.

IBM Tivoli Storage Productivity Center Standard Edition

Contains all the functions for data management, disk management, fabric management, and tape management, plus analytical functions:

- Configuration Analysis
- Configuration History
- SAN Planner
- Storage Optimizer

IBM Tivoli Storage Productivity Center Select

Contains all the functions for data management, disk management, fabric management, and tape management, plus the following analytical functions:

- Configuration Analysis
- Configuration History
- SAN Planner
- Storage Optimizer

These analytical functions are available for the following storage systems:

- System Storage DS3000
- System Storage DS4000
- System Storage DS5000
- XIV system
- SAN Volume Controller
- Storwize V7000

IBM Tivoli Storage Productivity Center for Data Edition

Includes data management, basic tape, disk, and fabric management, database monitoring, and chargeback functions, but does not include performance monitoring functions.

IBM Tivoli Storage Productivity Center for Disk Edition

Includes basic disk, fabric, tape, and data management functions and storage system performance monitoring, but does not include fabric performance monitoring, chargeback, and database monitoring functions.

IBM Tivoli Storage Productivity Center for Disk Select

Provides storage device configuration, performance monitoring, and management of Storage Area Network (SAN) attached devices from a single console.

This product provides basic disk functions and performance monitoring capabilities for the following storage systems:

- System Storage DS3000
- System Storage DS4000
- System Storage DS5000
- XIV system
- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

IBM Tivoli Storage Productivity Center Basic Edition

Includes basic disk, fabric, tape, and data management functions, but does not include chargeback, database monitoring, and performance monitoring functions.

Use the following tables to view the specific functions, reports, and elements of the user interface that are available in each license.

Table 3. Available functions in each Tivoli Storage Productivity Center license

Function	Basic Edition	Disk Edition	Disk Select Edition	Data Edition	Standard Edition	Standard Select Edition
Administrative Services						
Device discovery	Yes ¹	Yes ¹	Yes ¹	Yes	Yes	Yes
Data Sources	Yes ²	Yes ²	Yes ²	Yes ⁷	Yes	Yes
Configuration	Yes ³	Yes ³	Yes ³	Yes	Yes	Yes
IBM Tivoli Storage Productivity Center						
Configuration Utility	Yes	Yes	Yes	Yes	Yes	Yes
Configure Devices	Yes	Yes	Yes	Yes	Yes	Yes
Job Management	Yes	Yes	Yes	Yes	Yes	Yes
Reporting	Yes	Yes	Yes	Yes	Yes	Yes

Table 3. Available functions in each Tivoli Storage Productivity Center license (continued)

Function	Basic Edition	Disk Edition	Disk Select Edition	Data Edition	Standard Edition	Standard Select Edition
<ul style="list-style-type: none"> • Batch Reports • User ID reports • System Reports • Rollup Reports • Data Source Reports 	Not all reports are available in the Basic edition. See Table 4 on page 19 for a list of available reports.	Not all reports are available in the Disk edition. See Table 4 on page 19 for a list of available reports.	Not all reports are available in the Select edition. See Table 4 on page 19 for a list of available reports.	Not all reports are available in the Data edition. See Table 4 on page 19 for a list of available reports.	See Table 4 on page 19 for a list of available reports.	See Table 4 on page 19 for a list of available reports.
Topology Viewer	Yes	Yes	Yes	Yes	Yes	Yes
Monitoring						
<ul style="list-style-type: none"> • Probes • TPC Server Probes 	<ul style="list-style-type: none"> • Yes⁴ • No 	<ul style="list-style-type: none"> • Yes⁴ • No 	<ul style="list-style-type: none"> • Yes⁴ • No 	<ul style="list-style-type: none"> • Yes • No 	<ul style="list-style-type: none"> • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes
Storage Resource Group Management	Yes	Yes	Yes	Yes	Yes	Yes
Analytics						
<ul style="list-style-type: none"> • Configuration History • Storage Optimizer • SAN Planner • Configuration Analysis 	<ul style="list-style-type: none"> • No • No • No • No 	<ul style="list-style-type: none"> • No • No • No • No 	<ul style="list-style-type: none"> • No • No • No • No 	<ul style="list-style-type: none"> • No • No • No • No 	<ul style="list-style-type: none"> • Yes • Yes • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes • Yes • Yes
Alerting						
<ul style="list-style-type: none"> • Alert Log • Authentication Configuration Alerts 	<ul style="list-style-type: none"> • Yes⁵ • Yes 	<ul style="list-style-type: none"> • Yes⁵ • Yes 	<ul style="list-style-type: none"> • Yes⁵ • Yes 	<ul style="list-style-type: none"> • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes
Data Manager						
Monitoring						
<ul style="list-style-type: none"> • Groups • Pings • Scans • Profiles 	<ul style="list-style-type: none"> • Yes • Yes • No • No 	<ul style="list-style-type: none"> • Yes • Yes • No • No 	<ul style="list-style-type: none"> • Yes • Yes • No • No 	<ul style="list-style-type: none"> • Yes • Yes • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes • Yes • Yes
Alerting	Yes ⁶	Yes ⁶	Yes ⁶	Yes	Yes	Yes
Policy Management	No	No	No	Yes	Yes	Yes

Table 3. Available functions in each Tivoli Storage Productivity Center license (continued)

Function	Basic Edition	Disk Edition	Disk Select Edition	Data Edition	Standard Edition	Standard Select Edition
Reporting	Yes Not all reports are available in the Basic edition. See Table 4 on page 19 for a list of available reports.	Yes Not all reports are available in the Disk edition. See Table 4 on page 19 for a list of available reports.	Yes Not all reports are available in the Select edition. See Table 4 on page 19 for a list of available reports.	Yes	Yes	Yes
Data Manager for Databases	No	No	No	Yes	Yes	Yes
Data Manager for Chargeback	No	No	No	Yes	Yes	Yes
Disk Manager						
Storage Subsystems	Yes	Yes	Yes	Yes	Yes	Yes
<ul style="list-style-type: none"> • Create Volumes • Create Virtual Disks • Launch Element Manager 	<ul style="list-style-type: none"> • Yes • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes • Yes
Storage Optimizer	No	No	No	No	Yes	Yes
SAN Planner	No	No	No	No	Yes	Yes
Monitoring						
<ul style="list-style-type: none"> • Groups • Jobs • Storage Subsystem Performance 	<ul style="list-style-type: none"> • Yes • Yes • No 	<ul style="list-style-type: none"> • Yes • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes • No 	<ul style="list-style-type: none"> • Yes • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes • Yes
Alerting	Yes	Yes	Yes	Yes	Yes	Yes
Profile Management	No	Yes	Yes	No	Yes	Yes
Reporting						

Table 3. Available functions in each Tivoli Storage Productivity Center license (continued)

Function	Basic Edition	Disk Edition	Disk Select Edition	Data Edition	Standard Edition	Standard Select Edition
<ul style="list-style-type: none"> • Groups • Storage Subsystems • Storage Subsystem Performance 	<ul style="list-style-type: none"> • Yes • Yes • No 	<ul style="list-style-type: none"> • Yes • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes • No 	<ul style="list-style-type: none"> • Yes • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes • Yes
Fabric Manager						
Fabrics	Yes	Yes	Yes	Yes	Yes	Yes
Monitoring						
<ul style="list-style-type: none"> • Groups • Jobs • Switch Performance Monitors 	<ul style="list-style-type: none"> • Yes • Yes • No 	<ul style="list-style-type: none"> • Yes • Yes • No 	<ul style="list-style-type: none"> • Yes • Yes • No 	<ul style="list-style-type: none"> • Yes • Yes • No 	<ul style="list-style-type: none"> • Yes • Yes • Yes 	<ul style="list-style-type: none"> • Yes • Yes • Yes
Alerting	Yes	Yes	Yes	Yes	Yes	Yes
Reporting: Switch Performance	No	No	No	No	Yes	Yes
Tape Manager	Yes	Yes	Yes	Yes	Yes	Yes
Element Manager	Yes	Yes	Yes	Yes	Yes	Yes
Replication Manager	Yes	Yes	Yes	Yes	Yes	Yes

Table 3. Available functions in each Tivoli Storage Productivity Center license (continued)

Function	Basic Edition	Disk Edition	Disk Select Edition	Data Edition	Standard Edition	Standard Select Edition
Notes:						
1. You cannot run the following discovery jobs in the Basic, Disk, and Disk Select editions: Netware Filer, Windows Domain, NAS, SAN FS, and VMware VI Data Source.						
2. You cannot associate VMware VI Data Sources or TPC Servers with Tivoli Storage Productivity Center in the Basic, Disk, and Disk Select editions.						
3. The following nodes under Configuration are not available in the Basic, Disk, and Disk Select editions.						
<ul style="list-style-type: none"> • License Keys • Quota and Constraint email Address Rules • Scan/Probe Agent Administration • Manual NAS/Netware Server Entry • NetWare Tree Logins • Resource History Retention for Databases • Removed Resource Retention for Databases • Configuration History Settings (also not available in the Data edition) 						
4. Use probes and Storage Resource agents to collect asset and capacity information about storage entities in the Disk, Basic, and Disk Select editions. To collect specific file level information (for example, usage statistics), you must use the Data or Standard editions to run scans on your file systems. Note: You cannot probe hypervisors in the Basic, Disk, and Select editions.						
5. The following nodes under Alert Log are not available in the Basic, Disk, and Disk Select editions:						
<ul style="list-style-type: none"> • Directory • User • OS User Group • Configuration Analysis (also not available in the Data edition) • Hypervisor 						
6. The following nodes under Data Manager > Alerting are not available in the Basic, Disk, and Disk Select editions:						
<ul style="list-style-type: none"> • Directory Alerts • Hypervisor Alerts • Other NAS Alerts 						
7. Data Sources for Data Edition does not support TPC Server Probes.						

Table 4. Available reports in each IBM Tivoli Storage Productivity Center license

Reports	Basic Edition	Disk Edition	Disk Select Edition	Data Edition	Standard Edition	Standard Select Edition
Batch reports	No	Yes ^{1,2}	Yes ^{1,2}	Yes ²	Yes	Yes
System Reports:						
• Data	• Yes ⁵	• Yes ⁵	• Yes ⁵	• Yes	• Yes	• Yes
• Fabric	• Yes ³	• Yes ³	• Yes ³	• Yes ³	• Yes	• Yes
• Disk	• No	• Yes	• Yes	• No	• Yes	• Yes
Rollup Reports	No	No	No	No	Yes	Yes
Data Source Reports	Yes	Yes	Yes	Yes	Yes	Yes

Table 4. Available reports in each IBM Tivoli Storage Productivity Center license (continued)

Reports	Basic Edition	Disk Edition	Disk Select Edition	Data Edition	Standard Edition	Standard Select Edition
Asset reports:						
• By Cluster	• Yes	• Yes	• Yes	• Yes	• Yes	• Yes
• By Computer	• Yes	• Yes	• Yes	• Yes	• Yes	• Yes
• By Hypervisor	• No	• No	• No	• Yes	• Yes	• Yes
• By Storwize V7000U	• No	• No	• No	• Yes	• Yes	• Yes
• By OS Type	• Yes	• Yes	• Yes	• Yes	• Yes	• Yes
• By Storage Subsystem	• Yes	• Yes	• Yes	• Yes	• Yes	• Yes
• System-wide	• Yes ⁴	• Yes ⁴	• Yes ⁴	• Yes	• Yes	• Yes
Availability reports	Yes	Yes	Yes	Yes	Yes	Yes
TPC-wide Storage Space reports:						
• Disk Space	• Yes	• Yes	• Yes	• Yes	• Yes	• Yes
• File System Space	• Yes	• Yes	• Yes	• Yes	• Yes	• Yes
• Consumed File System Space	• Yes	• Yes	• Yes	• Yes	• Yes	• Yes
• Available File System Space	• Yes	• Yes	• Yes	• Yes	• Yes	• Yes
Usage reports	No	No	No	Yes	Yes	Yes
Usage Violation reports	No	No	No	Yes	Yes	Yes
Backup reports	No	No	No	Yes	Yes	Yes
Monitored Computer Storage Space reports	Yes	Yes	Yes	Yes	Yes	Yes
Storage Subsystem reports	Yes	Yes	Yes	Yes	Yes	Yes
Storage Subsystem Performance reports	No	Yes	Yes	No	Yes	Yes
Switch Performance reports	No	No	No	No	Yes	Yes
Tape Library report	Yes	Yes	Yes	Yes	Yes	Yes

Table 4. Available reports in each IBM Tivoli Storage Productivity Center license (continued)

Reports	Basic Edition	Disk Edition	Disk Select Edition	Data Edition	Standard Edition	Standard Select Edition
Notes:						
1. You must deploy Storage Resource agents to enable batch reports for these licenses. A Storage Resource agent must be deployed on the computer where you want to run batch reports.						
2. The following batch reports are not available in the following editions:						
<ul style="list-style-type: none"> • Basic edition: No batch reports are available • Disk edition: Rollup Reports, Usage, Usage Violations, Backup, Switch Performance • Disk Select edition: Rollup Reports, Usage, Usage Violations, Backup, Switch Performance • Data edition: Storage Subsystem Performance, Switch Performance 						
3. The following Fabric system reports are not available in the Basic, Disk, Disk Select, and Data editions: Top switch Ports Data Rate Performance, Top Switch Ports Packet Rate Performance, Switch Port Errors, Switch Performance.						
4. The following Asset > System-wide reports are not available in the Basic, Disk, and Disk Select editions: Unmanaged Virtual Machines, Virtual Machines with no VMWare Agent, Monitored Directories, Unmanaged Computers, Users, OS User Groups.						
5. The following Data system reports are available with Basic Edition, Disk, and Disk Select Edition:						
<ul style="list-style-type: none"> • Disk Space Summary • Disk Defects • Computer Storage Availability • Computer Disk Space • Available File System Space 						

If you have purchased a license for IBM Tivoli Storage Productivity Center, Table 5 lists the products for which you can purchase a license upgrade for your system.

Table 5. License upgrade scenarios for IBM Tivoli Storage Productivity Center

IBM Tivoli Storage Productivity Center license	First license upgrade	Second license upgrade
Basic Edition	Disk Edition	Standard Edition
Basic Edition	Data Edition	Standard Edition
Basic Edition	Disk Select Edition	Standard Edition
Basic Edition	Standard Edition	
Disk Select Edition	Disk Edition	Standard Edition
Disk Select Edition	Data Edition	Standard Edition
Disk Select Edition	Standard Select Edition	Standard Edition
Disk Select Edition	Standard Edition	
Data Edition	Disk Select Edition	Standard Edition
Data Edition	Standard Edition	
Disk Edition	Disk Select Edition	Standard Edition
Disk Edition	Standard Edition	
Standard Select Edition	Standard Edition	
Standard Edition		

Data Manager

Data Manager is a comprehensive file and capacity management solution for heterogeneous storage environments. Data Manager includes enterprise-wide reporting and monitoring, policy-based management and automated capacity provisioning for Direct Attached Storage (DAS), network attached storage (NAS), and SAN environments.

Data Manager helps you improve storage utilization, plan for future capacity, and ensure availability by providing storage on demand for file systems. Use Data Manager to perform the following functions:

- Discover and monitor disks, partitions, shared directories, and servers.
- Monitor and report on capacity and utilization across platforms to help you to identify trends and prevent problems.
- Provides a wide variety of standardized reports about filesystems and storage infrastructure to track usage and availability.
- Provide file analysis across platforms to help you to identify and reclaim space used by non-essential files.
- Provide policy-based management and automated capacity provisioning for filesystems when user-defined thresholds are reached.

Using these functions, Data Manager helps you lower storage costs by:

- Improving storage utilization
- Enabling intelligent capacity planning
- Helping you manage more storage with the same staff
- Supporting application availability through computer uptime reporting and application database monitoring.

Information collected by Data Manager helps you understand what is really going on with data on your servers and in your storage environment. View when files are created, accessed, and modified, and by what group or user. This type of information helps system administrators map storage resources to the consumers of the resource. The ability to map storage consumption to storage hardware has become increasingly important as the size of open systems environments has increased.

In addition to understanding the current consumption and usage of data within the enterprise, Data Manager tracks the information over time. Not only does this historical view of storage consumption and utilization show usage trends over time, the system administrator can also see a projected use of storage in the future. System administrators can prepare for the need to purchase additional capacity in a planned proactive manner rather than reacting to out-of-space emergencies.

Use Data Manager policy functions to help you evaluate and control the usage and status of your enterprise storage management assets. Because Data Manager is policy-based, it has autonomic self-healing capabilities that can detect potential problems and automatically make adjustments based on the policies and actions you have established. Use the capability to provision storage based upon storage policies to expand a file system, and to allocate storage to a volume.

For example, Data Manager can notify you when your system is running out of disk space or warn you of impending storage hardware failure. By alerting you to

these and other issues related to your stored data, Data Manager enables you to prevent unnecessary system and application downtime.

Data Manager for Databases

Data Manager for Databases provides a set of policy-driven automated tools for managing storage capacity, availability, events, performance and assets in your relational databases. It can help you identify, evaluate, control and predict the storage needs of Relational Database Management Systems (RDBMSs), which include Oracle, Sybase SQL Server, Microsoft SQL Server, and UDB/DB2.

Data Manager for Databases is policy-based and through its autonomic self-healing capabilities it can detect potential problems and automatically make adjustments based on the policies and actions you have established. For example, it can notify you when your database tables are running out of storage space or warn you of a dropped tablespace. By alerting you to these and other issues related to your stored data, it enables you to prevent unnecessary system and application downtime.

The following table lists the RDBMS objects whose storage usage is monitored by Data Manager for Databases.

Oracle	SQL Server	Sybase	DB2/UDB
Instances	Instances	Instances	Instances
Databases	Databases	Devices	Databases
Tablespaces	Datafiles	Databases	Tablespaces
Datafiles	Tables	Fragments	Containers
Tables	Indexes	Tables	Tables
Indexes	Log Files	Indexes	Indexes
Extents			
Segments			
Redo Log files			
Control Files			
Archive Log			
Directories			

Note: Data Manager for Databases does not currently support the monitoring of clustered database applications.

Data Manager for Databases collects information about RDBMS storage assets and performs tasks against the RDBMS storage assets. Methods for managing and automating capacity utilization of your databases help you:

- Monitor storage assets associated with enterprise-wide databases and notification of potential problems before they occur
- Make intelligent capacity management decisions based on current and trended historical data
- See how much storage is being consumed by users, groups of users and OS's within the database
- Create policy-based management for databases when user-defined threshold is reached
- View storage utilization management from a database and application perspective

Data Manager for Chargeback

Data Manager for Chargeback uses the storage usage information gathered by Data Manager and Data Manager for Databases to generate invoices that chargeback for storage usage.

With Data Manager for Chargeback you can collect usage information on a departmental, group, or user level. You can allocate costs by storage usage by user, disk capacity by computer, tablespace, or file system/physical device. You can create cost centers by creating user, computer, or tablespace groups, allowing organization to chargeback individuals or business units for their storage usage. By understanding the costs associated with existing and future storage usage, you can improve the use of that storage and reduce the costs associated with its maintenances and upgrades.

Using Data Manager, you can run monitoring jobs that gather detailed statistics about the usage and trending of the storage consumed by the users and user groups within your organization. Using Data Manager for Databases, you can run monitoring jobs that gather detailed statistics about the usage and trending of the storage consumed within RDBMS Instances by the users and user groups within your organization.

In addition to providing invoicing for storage usage, Data Manager for Chargeback also integrates with the chargeback systems already implemented in your environment. It provides you with a higher level, application-specific CIMS output format which can be imported into CIMS applications. With Data Manager for Chargeback, you can export chargeback data for direct import into CIMS, giving you the ability to integrate your storage usage data with other enterprise chargeback information and processes.

Disk Manager

Disk Manager helps you manage SANs and heterogeneous storage from a single console.

With Disk Manager you can manage network storage components such as:

- IBM System Storage SAN Volume Controller
- IBM Storwize V7000
- Storwize V7000 Unified
- IBM TotalStorage Enterprise Storage Server (TotalStorage Enterprise Storage Server)
- Tivoli Storage Productivity Center disk systems (DS4000, DS5000, DS6000, and DS8000 series)
- IBM XIV Storage System
- Other storage subsystems that support the SMI-S standards

For those devices that use the native interfaces (DS8000, SAN Volume Controller, Storwize V7000, and XIV) device discovery is performed when the device is added to Tivoli Storage Productivity Center.

For those devices that use the SMI-S standards, device discovery is done using the Service Location Protocol (SLP), as specified by SMI-S. Configuration of the discovered devices is possible in conjunction with CIM agents associated with those devices, using the standard mechanisms defined in SMI-S. Disk Manager also gathers events, and can launch an element manager specific to each device.

Disk Manager can:

- Collect and store performance data and provides alerts
- Provide graphical performance reports
- Help optimize storage allocation
- Provide volume contention analysis

Through the use of data collection, setting of thresholds and use of performance reports, performance can be monitored for the TotalStorage Enterprise Storage Server, DS4000, DS5000, DS6000, DS8000, SAN Volume Controller, Storwize V7000, XIV, and any other storage subsystem that supports the SMI-S block server performance subprofile. The performance function starts with the data collection task, which captures performance statistics for the devices and stores the data in the database.

After data collection is done, you can set thresholds to identify exceptions for certain performance metrics, depending on the type of device. Threshold checking is performed during data collection, and when performance is outside the specified bounds, alerts can be generated.

After performance data has been collected, you can configure Disk Manager to present graphical or text reports on the historical performance behavior of specified devices, or of specific internal components of those devices. The performance reports provide information about past period performance metrics or current performance in graphical form.

Note: Tivoli Storage Productivity Center supports discovery of and reporting on both FB and CKD volumes (and pools and arrays), but can create only FB volumes. Tivoli Storage Productivity Center does not support the creation of CKD volumes.

Managing TagmaStore CIM agents

The TagmaStore CIM agents are provided by Hitachi Data Systems for the TagmaStore storage subsystem. The TagmaStore CIM agent collects information from the TagmaStore storage system.

IBM Tivoli Storage Productivity Center supports the Hitachi Data Systems TagmaStore CIM Agent. For information about the CIM agent level supported, see http://www.ibm.com/support/docview.wss?rs=40&q1=support+matrix&uid=swg21386446&loc=en_US&cs=utf-8&lang=en. Click the appropriate version under **Storage**. Tivoli Storage Productivity Center also supports the TagmaStore as a Storage Virtualizer. Tivoli Storage Productivity Center can display information for virtual disks and local disks.

Tivoli Storage Productivity Center cannot provide correlation information between the virtual storage used by TagmaStore and the volumes created on the storage due to an existing limitation of the CIM agent 5.8 from Hitachi Data Systems. However, Tivoli Storage Productivity Center reports correctly display the correlation between volumes created on the local storage and the local disks. This limitation has no impact on the topology but it does affect several reports that show the correlation:

- **Data Manager > Reporting > Asset > By Storage Subsystem > <HDS_device> > Disks > Volumes** (does not show the relation between the disk and Volume or volume in the tree)
- **Disk Manager > Reporting > Storage Subsystem > Volume to Backend Volume Assignment**

- **Disk Manager > Reporting > Storage Subsystem > Computer Views** (because Tivoli Storage Productivity Center cannot populate the information for volumes created on the virtual disks)

All volumes are created from a storage pool that is allocated from a primordial storage pool and an imported primordial pool. A volume cannot be created over both local and virtual extents.

For volume correlation, the host computer must have the Storage Resource agent deployed and the TagmaStore device must be in the same SAN fabric. There must also be a zone configured in the active zoneset between the ports of the host computer and the ports of the TagmaStore device. The Storage Resource agent must be configured for the fabric to which the host is connected.

For back-end correlation, the TagmaStore device ports and back-end subsystem ports must be in the same zone and the back-end subsystem has assigned storage volumes to all ports of the TagmaStore device.

Tape Manager

IBM Tivoli Storage Productivity Center Tape Manager helps you manage your SMI-S compliant tape libraries.

Tape Manager lets you manage those tape libraries which are based on SMI-S 1.1 profile for tape libraries.

You must have at least Tape Operator authority to access the functions in the Tape Manager node. To view your authority level, open the "Role to Group Mappings" window in the **Administrative Services > Configuration** node in the navigation tree.

With Tape Manager you can:

- Discover SMI-S compliant tape libraries attached to your SAN fabrics
- View tape library information
- Monitor tape libraries by probing for data and viewing alerts
- Create and view asset and capacity reports using data collected from your tape libraries
- Use the topology viewer to view information about tape libraries

Fabric Manager

Fabric Manager helps you manage the SAN fabric that connects the host systems and applications to the storage devices. It is a comprehensive management solution for multi-vendor SANs and includes automatic resource and topology discovery, monitoring and alerts, and zone control.

Fabric Manager is an enterprise scalable solution architected to ANSI SAN standards, allowing you to choose the products right for your environment. It helps you:

- Simplify the task of SAN management and configuration
- Ensure SAN availability
- Improve SAN return on investment

Fabric Manager is able to do SAN fabric performance and capacity management reporting and monitoring. Zoning is one of many SAN security options. Zoning is

the only configuration supported for this release. Along with subsystem based volume masking, it is almost universally employed to ensure that only systems that are authorized can access selected data. Zoning is provided by the switch at the port level, so that, for example, a host on port X can access the subsystem connected by port Y. When business needs change, often zones must change also.

Zone control is enabled from a centralized location. Existing zones can be discovered and the members that make up the zone can be viewed and modified by adding or deleting them. New zones can be created and existing zones can be deleted. Fabric Manager makes working with zones easy by providing a GUI that discovers SAN devices and makes them available to add or remove from zones as individual devices.

The ability to do switch performance and capacity management reporting and monitoring can help you to determine if more bandwidth is needed. Wide-area and local-area IP networks, and storage area networks (SANs) move data from one place to another. The management of the bandwidth is needed to continually monitor for link utilization and errors. Data needs to be gathered to tune resources, balance workloads, and do capacity planning.

With Fabric Manager, the comprehensive management of multi-vendor SANs can help simplify the IT infrastructure by providing automatic resource and topology discovery, monitoring, and alerts, and zone control. It brings all the sources of information about SAN topology and configuration into a single place, and creates topology mapping of the SANs. This topology offers both host-centric and device-centric views that can be displayed graphically. The SAN topology display tracks all topology and configuration changes through in-band, out-of-band, and SMI-S monitoring. Without this type of centralized topology, information from a number of sources, such as Element Management tools, device logs, and SNMP traps would have to be continually monitored and manually correlated to determine the current SAN configuration and topology.

Fabric Manager provides you with the ability to view events happening in your SAN environment and records state changes. The events are displayed in a color-coded fashion and can be further customized to reflect organizational priorities. It will forward events signaling topology changes or updates to the IBM Tivoli Enterprise Console, another SNMP manager, or both. Source side filtering by the Fabric Manager enables you to control what events get sent to your centralized console. Filtering helps ensure that the administrators get the information they need and are not overwhelmed by extraneous information.

Graphical user interface

Use the navigation tree on the left side of the user interface to access the functions in Tivoli Storage Productivity Center.

Figure 1 shows the main window, which is displayed when you log in to Tivoli Storage Productivity Center.

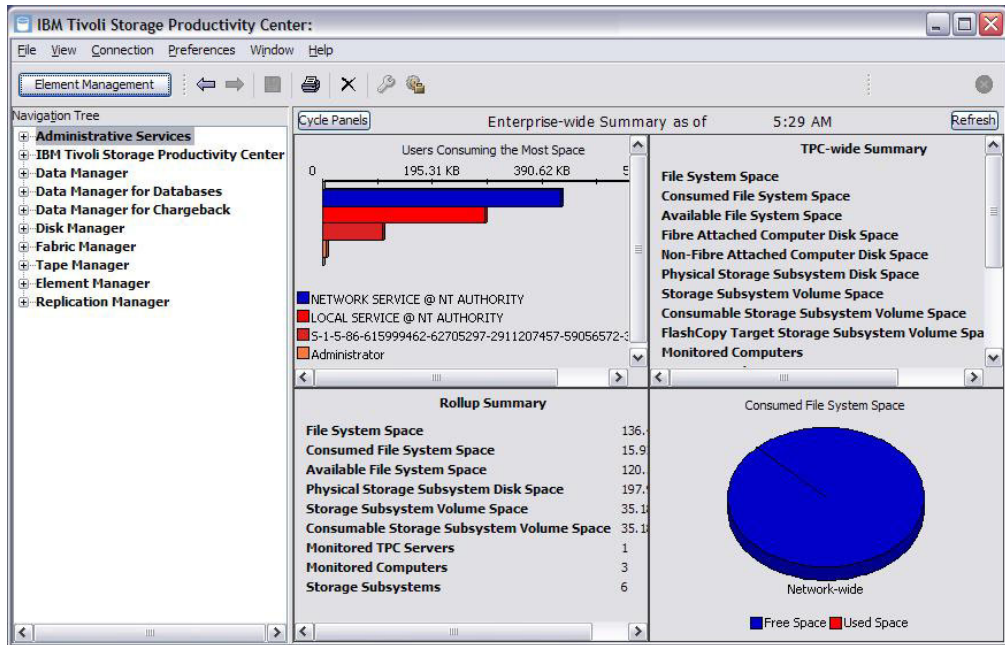


Figure 2. Tivoli Storage Productivity Center Main Window

The Tivoli Storage Productivity Center user interface consists of a navigation tree and a content pane.

- **Navigation tree**



The navigation tree is on the left side of the main window. You can expand, select, and collapse the tree nodes to navigate through the interface.

When you start Tivoli Storage Productivity Center, the navigation tree is expanded to show all the high-level functions. You can drill down on an element in the tree by clicking it or by clicking the expand icon.



When you right-click a node a pop-up menu is displayed. Use this menu to perform additional actions for the node.

Expanding and collapsing the navigation tree: If the navigation tree gets too large, or if you want to return it to its original state, right-click the major nodes of the tree and select **Collapse Branch** or **Expand Branch**. If you right-click the **Tivoli Storage Productivity Center** node and select **Collapse Branch**, that entire subtree collapses. Then, right-click the main **Tivoli Storage Productivity Center** node and select **Expand Branch** to return the navigation tree to its original state—expanded only to show the main functions.

Printing the navigation tree: Right-click a node and choose **Print Branch** to print all the subnodes within that node.

- **Content pane.** The content pane is on the right side of the main window. When you select a node in the navigation tree, a corresponding panel is displayed in the content pane. Use panels in the content pane to define and run the different functions (for example, data collection schedules, alerts, reports, and so on) that are available within Tivoli Storage Productivity Center.

The GUI works like a Web browser, whether you are accessing it through the Web

or from the installed interface application. The  and  buttons on the tool bar move control back and forth through Tivoli Storage Productivity Center panels that you have been working on during a product session.

Navigation tree

The navigation tree shows all of the available resource management tasks, organized by major topic nodes for Administrative Services, IBM Tivoli Storage Productivity Center, Data Manager, Disk Manager, Fabric Manager, and Tape Manager.

Administrative Services

Provides product configuration functions for Tivoli Storage Productivity Center.

This node is always present.

Tivoli Storage Productivity Center

Provides operational functions for Tivoli Storage Productivity Center.

This node is always present.

Data Manager

Provides operational functions for the Data Manager.

This node is present when Data Manager is installed and licensed.

For Data Manager functions, you must use Data agents or Storage Resource agents to collect data about your storage resources. For example, use scans, pings, and probes to collect data about the computers on which the agents are installed.

Data Manager for Databases

Provides operational functions for Data Manager for Databases.

This node is present when Data Manager is installed and licensed.

Data Manager for Chargeback

Provides operational functions for Data Manager for Chargeback.

This node is present when Data Manager is installed and licensed.

Disk Manager

Provides operational functions for the Disk Manager.

This node is present when Disk Manager is installed and licensed. A subset of the functions is provided when you install the Data Server.

For Disk Manager functions, the subsystem CIMOMs and data collection are required (for example, through probes), except for devices that use Native API connection: IBM System Storage DS8000, IBM System Storage SAN Volume Controller, and IBM XIV Storage System.

Fabric Manager

Provides operational functions for the Fabric Manager.

This node is present when Fabric Manager is installed.

For Fabric Manager functions, the available agent types are CIM agents, in-band fabric agents, out-of-band fabric agents, and Storage Resource Agents. You must run a data collection job (for example, through probe jobs). If you want to perform zone control actions, the in-band fabric agents for QLogic and Cisco fabrics must be deployed. For Fabric Manager

features such as zone control and performance monitoring, you must perform data collection from performance monitors. For a list of current agent types supported for the Switch Performance Management functionality and the Fabric Zone Configuration functionality, see “Fabric Manager” on page 1147.

Tape Manager

Provides operational functions for the tape library.

This node is present when the Device server is installed.

For Tape Manager functions, the library CIMOMs and data collection are required (for example, probes).

Element Manager

Allows you to access the **Element Manager** tab of the Configuration Utility where you can configure and administer IBM System Storage DS8000 and other storage subsystems.

Replication Manager

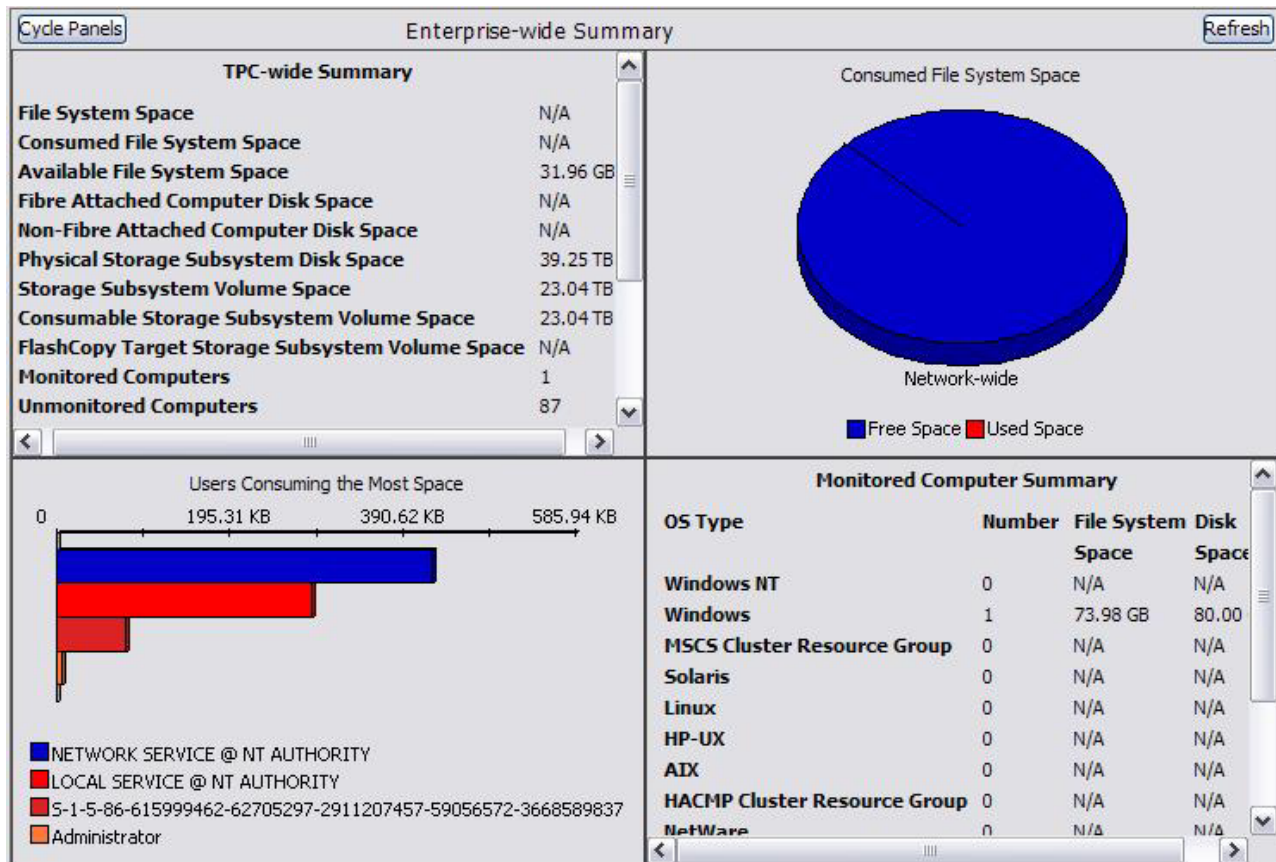
Allows you to access the **Replication Manager** tab of the Configuration Utility where you can manage IBM Tivoli Storage Productivity Center for Replication and to set the triggering condition and actions for alerts associated with Tivoli Storage Productivity Center for Replication.

Dashboard

The dashboard provides a concise, yet detailed overview of the health of your storage environment. It quickly points out potential problem areas that need further investigation.

When you first log in to IBM Tivoli Storage Productivity Center you are asked if you want to collect statistics. You must answer yes in order for the dashboard information to be displayed. The dashboard then appears whenever you start the product or close all the active windows in a running session.

The dashboard contains four display areas and seven panels that you can cycle through to view your environment's storage information.



Using the dashboard:

Use the dashboard to monitor the overall status of your storage resources.

Purpose

When you first log on to Tivoli Storage Productivity Center or close all the active panels, the topic pane shows the "dashboard". The dashboard provides a concise, yet detailed overview of the status of your storage environment. Use it to quickly point out potential problem areas that need further investigation. The dashboard contains four displayable areas and eight panels that you can cycle through to view the storage information that has been collected about your environment.

Keep in mind the following considerations when viewing the information about the dashboard:

- The information in the dashboard is collected by data collection schedules and aggregated by the History Aggregator. To ensure that the latest information is displayed in the dashboard, run the history aggregator regularly. See the information center for instructions about how to create and schedule history aggregator jobs.
- The value in the **All** field on the Alerts Pending panel might be higher than the sum of the alerts shown in the other fields on that panel. This situation occurs because some alerts are not associated with a specific category, and are reflected in the value that is shown in the **All** field only.
- The names of some fields have been updated. For your reference, the original names of those fields are shown within parentheses () next to their new names in this topic.

- Click **Refresh** to update the displayed panels. You must click **Refresh** to ensure that the latest information from the history aggregator is displayed in the dashboard.
- Click **Cycle Panels** to cycle through the panels in the dashboard and view information about your storage assets.

TPC-wide Summary panel

Use this panel to view an enterprise-wide summary of your storage through the following values:

File System Space (Capacity)

Total amount of unique file system storage space on computers where a Data agent or Storage Resource agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this field:

- File systems that become missing after a probe.
- File systems discovered on virtual machines where an agent is installed.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Consumed File System Space (Filesystem Used Space)

Total amount of unique file system storage space that is used or consumed by the computers where a Data agent or Storage Resource agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this field:

- File systems that become missing after a probe.
- File systems on discovered hypervisors that are assigned to virtual machines where an agent is installed.
- File systems on virtual machines where an agent is installed, but the hypervisors for the virtual machines have not been discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Available File System Space (Filesystem Free Space)

Total amount of unique file system storage space that is not used or available to the computers where a Data agent or Storage Resource agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this field:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.

Fibre Attached Computer Disk Space (Computer Fibre Attached Disk Space)

Total amount of unique disk space on computers where a Data agent or Storage Resource agent is installed and the disks are attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this field:

- Computer disks that are behind a fibre channel port that has not been identified by an agent.

- Computer disks that are not behind a fibre channel port.
- Computer disks discovered on virtual machines where an agent is installed.
- Computer disks that become missing after a probe.

Non-Fibre Attached Computer Disk Space (Computer Non-fibre Attached Disk Space) Total amount of unique disk space on computers where a Data agent or Storage Resource agent is installed and the disks are not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this field:

- Computer disks that are behind a fibre channel port.
- Computer disks discovered on virtual machine where an agent is installed.
- Computer disks that become missing after a probe.

Physical Storage Subsystem Disk Space (Storage Subsystem Physical Disk Space) Total amount of unique internal physical disk space on the storage subsystem arrays that are monitored by subordinate servers and the master server. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this field:

- Storage subsystem disks that become missing after a storage subsystem probe.
- Back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller) that are attached from a storage subsystem.
- Storage subsystem disks that are on storage subsystems that have not been probed by the subordinate servers or the master server.

Storage Subsystem Volume Space (LUN Capacity)

Total amount of unique storage subsystem volume space and virtual storage volume space on the storage subsystem arrays that are monitored by subordinate servers and the master server. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this field:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller).
- Storage subsystem disks that are on storage subsystems that have not been probed by the subordinate servers or the master server.

Consumable Storage Subsystem Volume Space (Usable LUN Capacity)

Amount of file system storage space for a file system. Total amount of unique storage subsystem volume space and virtual storage volume space on monitored storage subsystem arrays that can be assigned or are assigned to systems within the network. Tivoli Storage Productivity Center does not include the following columns in its calculation of the value for this field:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are used as FlashCopy targets.

- Storage subsystem volumes or virtual storage volumes that are identified as Business Continuity Volume extended (BCVx).
- Storage subsystem volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database.
- Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller).

FlashCopy Target Storage Subsystem Volume Space (FlashCopy Target Capacity)

Total amount of unique storage subsystem volume space and virtual storage volume space on monitored storage subsystem arrays that is identified as FlashCopy target space. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this field:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are not used as FlashCopy targets.

Monitored Computers (Servers)

Total number of computers where a Data agent or Storage Resource agent is installed and are detected by the Data server. Tivoli Storage Productivity Center does not include the following systems in its calculation of the value for this field:

- Systems where an agent is installed but was not started or was unable to communicate with the Data server.
- Systems where an agent is not installed.
- Systems where an agent is installed but that agent communicates with a different Tivoli Storage Productivity Center server.
- Systems that have been removed from the Administrative Services > Data Sources > Data/ or Storage Resource agents panel.
- Hypervisors

Unmonitored Computers (Servers)

Total number of computers found in a Windows domain that do not have a Data agent or Storage Resource agent communicating with the Data server. Tivoli Storage Productivity Center does not include the following systems in its calculation of the value for this field:

- Systems that are not running a Windows operating system.
- Systems that are not in a Windows domain.
- Systems where an agent is installed and is communicating with a Data server.

Storage subsystems

Total number of storage subsystems that have been discovered and probed. Tivoli Storage Productivity Center does not include the following storage subsystems in its calculation of the value for this field:

- Storage Subsystems that have been discovered but not probed.
- Storage Subsystem that have not been discovered.
- Storage Subsystems that have been removed from the Disk Manager > Storage Subsystems panel.

Users Total number of user accounts discovered by Data agents or Storage Resource agents.

Disks Total number of unique disks discovered by Data agent or Storage Resource agent probes, hypervisor probes, network storage filer probes, and storage subsystem probes. This value does not include disks that have become missing after a successful probe.

Storage Subsystem Volumes (LUNS)

Total number of unique storage volumes and virtual storage volumes discovered by storage subsystem probes. This value does not include storage volumes or virtual storage volumes that have become missing after a successful storage subsystem probe.

Filesystems

Total number of unique file systems discovered by probes. This value excludes file systems that have become missing after a successful probe.

Directories

Total number of directories discovered by scans. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this field:

- Directories within file systems that were discovered by a probe but were not scanned.
- Directories that have become missing after a successful probe.

Files

Total number of files discovered by scans. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this field:

- Files within file systems that were discovered by a probe but were not scanned.
- Files that have become missing after a successful probe.

Rollup Summary panel

Use this panel to view a summary of information collected by the master server and by Tivoli Storage Productivity Center servers that are subordinate servers.

File System Space (Capacity)

Total amount of unique file system storage space on the computers that are monitored by subordinate servers and the master server. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this field:

- File systems that become missing after a probe.
- File systems discovered on virtual machines where the Data agent or Storage Resource agent is installed.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.
- File systems that have not been probed by subordinate servers or the master server.

Consumed File System Space (Used Space)

Total amount of unique file system storage space that is used or consumed by the computers that are monitored by subordinate servers and the master server. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this field:

- File systems that become missing after a probe.

- File systems on discovered hypervisors that are assigned to virtual machines where the Data agent or Storage Resource agent is installed.
- File systems on virtual machines where the Data agent or Storage Resource agent is installed, but the hypervisors for the virtual machines have not been discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.
- File systems that have not been probed by subordinate servers or the master server.

Available File System Space (Filesystem Free Space)

Total amount of unique file system storage space that is not used or available to the systems that are monitored by subordinate servers and the master server. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this field:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network-attached storage server.
- File systems that are not mounted.
- File systems that have not been probed by subordinate servers or the master server.

Physical Storage Subsystem Disk Space (Storage Subsystem Physical Disk Space)

Total amount of unique internal physical disk space on monitored storage subsystem arrays. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this field:

- Storage subsystem disks that become missing after a storage subsystem probe.
- Back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller) that are attached from a storage subsystem.

Storage Subsystem Volume Space (LUN Capacity)

Total amount of unique storage subsystem volume space and virtual storage volume space on monitored storage subsystem arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this field:

- Storage subsystem volumes or virtual storage volumes that become missing storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller).

Consumable Storage Subsystem Volume Space (Usable LUN Capacity)

Total amount of unique storage subsystem volume space and virtual storage volume space on storage subsystem arrays that are monitored by subordinate servers and the master server and can be assigned or are assigned to systems within the network. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this field:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.

- Storage subsystem volumes or virtual storage volumes that are used as FlashCopy targets.
- Storage subsystem volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
- Storage subsystem volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database.
- Storage subsystem volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage subsystem virtualizers (for example, SAN Volume Controller).
- Storage subsystem disks that are on storage subsystems that have not been probed by the subordinate servers or the master server.

Monitored TPC Servers

Total number of subordinate servers that are monitored by the master server.

Monitored Computers (Servers)

Total number of computers that are monitored by Tivoli Storage Productivity Center and are detected by the subordinate and the master server. Tivoli Storage Productivity Center does not include the following systems in its calculation of the value for this field:

- Systems where an agent is installed but was not started or was unable to communicate with the Data server.
- Systems where an agent is not installed.
- Systems where an agent is installed but that agent communicates with a Tivoli Storage Productivity Center server that is not a subordinate server or the master server.
- Systems that have been removed from the Administrative Services > Data Sources > Data/ or Storage Resource agents panel.

Storage subsystems

Total number of storage subsystems that have been discovered and probed by subordinate servers and the master server. Tivoli Storage Productivity Center does not include the following storage subsystems in its calculation of the value for this field:

- Storage Subsystems that have been discovered but not probed.
- Storage Subsystem that have not been discovered.
- Storage Subsystems that have been removed because of the storage subsystem settings on the Removed Resource Retention panel.

Consumed File System Space (Filesystem Used Space panel)

Use the chart on this panel to view the total amount of used file system storage space across the entire network.

Users Consuming the Most Space (chart) panel

Use this panel to view the following information:

- A list of the users consuming the most space within your environment
- A graphical view of the amount of space they are consuming

The chart on this panel shows information about users across the network. If multiple hosts contain the same user name, the value displayed for that user represents the space consumed on all the hosts. For example, if the user "root"

exists on three hosts that are monitored by Tivoli Storage Productivity Center, then the value displayed on this chart represents the space consumed by "root" on those three hosts.

Note: The space usage of administrative users within your environment is not included as part of this dashboard report.

Monitored Computer Summary panel

Use this panel to view the following information:

- The number of monitored computers running each Operating System type
- Total amount of file system and disk capacity per Operating System type

Filesystems with Least Free Space Percentage panel

Use this panel to view the monitored computers with the least percentage of free space, including the following values:

- Percent of free space on each computer
- Total file system capacity on each computer
- Name of the computers, file systems, and mount points

Users Consuming the Most Space Report panel

Use this panel to view the following information:

- A list of the users consuming the most space within your environment
- The total amount of storage each user is consuming

Alerts Pending panel

Use this panel to view the active alerts that are still pending.

Note: The number of alerts on this panel might be less than the number of alerts shown in the IBM Tivoli Storage Productivity Center > Alerting > Alert Log. Alerts for failed jobs (pings, probes, and scans) are not listed in separate rows on this panel, but are included in the total displayed in the **All** row.

These alerts are organized into the following categories:

- All
- Computer
- Disk
- Filesystem
- Directory
- User Quota
- OS User Group Quota
- Directed to *user*
- Storage Subsystem
- Switch
- Fabric
- Endpoint Device
- External
- Tape Library

- Hypervisor
- Authentication Configuration
- Replication
- Configuration Analysis
- NAS Node
- NAS File Set
- Export

Working with the navigation tree and content pane

This section describes some methods for working with the nodes and windows in the navigation tree and content pane.

Right-clicking and pop-up menus:

Right-click nodes in the navigation tree to access menus that contain actions you can perform.

As you navigate through the IBM Tivoli Storage Productivity Center user interface, notice that you can right-click nodes in the navigation tree to access menus that contain actions related to that node. Use these actions to quickly access many product functions.

The following actions are available for the nodes in the navigation tree:

Option	Description
Create <object>	Create an object, including a group, data collection schedule, policy, batch report, or alert.
Delete	Delete the selected object definition, such as a data collection schedule, policy, batch report, group, or alert. This action is not available for definitions owned by TPCUser.
Rename	Rename the object definition. This action is not available for definitions owned by TPCUser.
Refresh	Refresh the list of objects in the navigation tree. The navigation tree is not updated automatically. You must click Refresh for a node to view the latest objects related to that node. For example, right-click the Tivoli Storage Productivity Center > Monitoring > Probes node and select Refresh to update the probe schedules that are listed in the navigation tree.
Job History	View the list of jobs related to a data collection or batch report schedule on the Job Management panel.
Run Now	Run a data collection or batch report schedule immediately.
Collapse Branch	Minimize the navigation tree under the selected node.
Expand Branch	Open the navigation tree under the selected node.
Print Branch	Print the contents of the navigation tree under the selected node.

Selecting panels in the content pane:

The content pane appears on the right side of the user interface and displays the panels related to the actions that you want to perform.

When you click a node in the navigation tree, a panel for that node opens in the content pane. Each time you select a new node, the panel for the previously

selected node is not deleted. Instead, it moves into the background and is

accessible by clicking  and  on the toolbar. You can also select currently active panels from the **Window > Show Window List** in the menu bar.

When viewing the active panel listed in the **Window** menu, a > appears next to the current panel (the panel that currently occupies the content pane). Objects that you are currently editing (such as data collection schedules) are marked with a black arrow if they have no pending changes that have not been saved; a solid red arrow indicates objects with pending changes that have not been saved.

Menu bar

Use the menu bar at the top of the main window to access some of the main functions within IBM Tivoli Storage Productivity Center.

The following sections describe the options available from the menu bar.

File menu:

Use the options in the File menu to save objects and jobs, export report data, print, refresh alerts, close the current window, and exit the product session.

Menu Option	Description
Save	Save the changes for the current object (for example, groups, profiles, etc.) or job (for example, probes, scans, pings, etc.) that you are editing.
Save As	Save the changes for the current object or job that you are editing under a different name.
Export Data	Export the report data to other formats: <ul style="list-style-type: none">• Formatted Report file (.txt)• Comma delimited file (.csv)• Comma delimited file with headers• HTML file This allows you to import the data to another application, such as Excel, if needed.
Print	Print the current data in the content pane. You can print tabular reports or report charts. You can also save the output to a PDF file or HTML file, and print those files at a later time.
Print Tree	Print the expanded contents of the navigation tree. You can also save output to a PDF or HTML file, and print those files at a later time.
Refresh Alerts	Refresh all the alerts from the database repository and update the GUI by highlighting the appropriate objects and showing all the new alerts in the alert logs.
Close	Close the current window that occupies the content pane.
Exit	End the current session.

View menu:

Use the options in the View menu to hide or show the navigation tree and highlight the node in the navigation tree for the currently open window.

Menu Option	Description
Tree	Hide and show the navigation tree. If you want to increase the size of the content pane to occupy the entire window, clear the check box for the Tree option. To redisplay the navigation tree next to the content pane in the main window, select the Tree option.
Current Page in Tree	Display the highlighted node on the navigation tree that was responsible for displaying the data on the content pane. Use this option to reorient the navigation tree with the window that is currently displayed in the content pane.

Connection menu:



Use the option in the Connection menu to end the current session and establish a new connection to the same IBM Tivoli Storage Productivity Center server or to a different server.

Menu Option	Description
New Connection	Initiate a new connection to a server from the current GUI session. For example, use this option to connect to the same server using a different user ID, or to a completely different server.
	All the current changes that are unsaved in the open windows will be lost when you establish a new connection.

Preferences menu:

Use the options in the Preferences menu to change the appearance of the user interface, edit the way the product manages open panels, determine what alerts are shown upon login, select what tab appears first when generating reports, and specify what unit of measurement to use when displaying storage sizes in reports.

Menu Option	Description	
Look and Feel	Windows	Set the look and feel of the GUI to appear like Microsoft Windows GUI.
	CDE/Motif	Set the look and feel of the GUI to appear like the UNIX Motif interface.
	Metal	Set the look and feel of the GUI to appear like Metal, the cross-platform Java look and feel.

Menu Option	Panel Retention	Description
Edit General	Panel Retention	<p>The Panel recall depth value specifies the number of windows that are linked in the content pane for access via the</p>  <p>and</p>  <p>buttons on the toolbar. The default is 20.</p> <p>The Panels kept in memory field specifies the number of windows to store in memory, which helps to optimize the system performance. The default is 2.</p> <p>Note: Any open views in the topology viewer (for example, an L2:Storage Subsystem view, an L0:Computers view) are considered part of the topology viewer panel and are listed as tabs at the top of the topology viewer panel. If the value in this field is a low number, such as the default of 2, then opening additional non-topology panels can exceed the number of panels that are kept in memory. This can result in the topology viewer panel refreshing and any tabbed views being closed. To avoid this situation, change the value in this field to a higher number, such as 10.</p>
	On login, show	<p>Control which alerts you will view upon logging in to a user interface session.</p> <p>You can select:</p> <ul style="list-style-type: none"> • No Alerts • Alerts Directed to <Userid> • All Active Alerts
	Initial Reporting Tab to Display	<p>Control what tab will be displayed when you first generate and view a report. You can select:</p> <ul style="list-style-type: none"> • Selection tab • Report tab • Chart tab (if Available)

Menu Option	Description
Advanced Options	<p>Select the unit of measurement that should be used to reflect the storage sizes included in IBM Tivoli Storage Productivity Center reports. This setting affects the results in online, exported, and Batch reports. You can choose from the following options:</p> <ul style="list-style-type: none"> • As Appropriate: Choose this option to automatically determine the unit measurement (KB, MB, GB, TB) to display for storage sizes in a report. • Kilobytes (KB): Choose this option to display storage sizes in reports using the kilobyte unit of measurement. For example, if the storage capacity for a computer is 100 megabytes, it is shown as 102400.00 KB in a capacity report. • Megabytes (MB): Choose this option to display storage sizes in reports using the megabyte unit of measurement. For example, if the storage capacity for a computer is 1 gigabyte, it is shown as 1024 MB in a capacity report. • Gigabytes (GB): Choose this option to display storage sizes in reports using the gigabyte unit of measurement. For example, if the free space on a computer is 500 megabytes, it is shown as 0.49 GB in a free space report. • Terabytes (TB): Choose this option to display storage sizes in reports using the terabyte unit of measurement. For example, if the free space on a computer is 350 gigabytes, it would be shown as 0.34 TB in a free space report.

Window menu:

Use the options in the Windows menu to close all the open windows that do not have any changes, display the dashboard, and list all the open windows.

Menu Option	Description
Close all Windows Without Changes	Close all of the open windows that do not contain any modifications. Any windows that contain active changes to editable objects will not be closed.
Show Dashboard	Display the IBM Tivoli Storage Productivity Center dashboard.
Show Window List	<p>Lists names of all open windows active in the content pane. Select one of these options to view the corresponding window.</p> <p>Beside each window name is an icon that indicates the status of the window:</p> <ul style="list-style-type: none"> • no icon - No changes were made to the window • solid red circle - Changes were made, but have not been saved • black arrow (>) - Current window displayed. No changes were made • solid red arrow (>) - Current window displayed. Changes were made, but have not been saved

Help menu:

Use the options in the Help menu to access online help for the product, launch the information center Web page, launch the Supported Products Web pages, launch the demonstrations Web page, and display product information (such as the version number).

Menu Option	Description
Help Topics	Select this menu item to launch IBM Tivoli Storage Productivity Center's online Help system. The Help system will appear in its own window and display a table of contents from which you can access information about how to use the product.
Help for Displayed Panel	Select this menu item to launch the online Help system and display help information for the product window that is currently displayed. You can also press F1 to launch the help topic for the product window that you are viewing.
Launch Information Center	<p>Select this menu item to launch the information center for IBM Tivoli Storage Productivity Center Information Center in a Web browser. The IBM Tivoli Storage Productivity Center Information Center contains detailed instructions that explain how to install, upgrade, configure, maintain, and use the Tivoli Storage Productivity Center product.</p> <p>See "Web browser support" on page 45 for information about Web browser support for launching the IBM Tivoli Storage Productivity Center Information Center.</p>
Supported Products Data	<p>Select the links under this menu item to launch Web pages that contain additional documentation for Disk Manager, Fabric Manager, and Data Manager. This documentation includes matrixes that show the devices supported by each of the managers.</p> <ul style="list-style-type: none">• IBM Tivoli Storage Productivity Center for Disk Supported Products. Click this item to view the Support for Tivoli Storage Productivity Center for Disk Web page. To view the devices supported by Disk Manager, scroll to the Integration/Operation section and click the link that matches your version of Tivoli Storage Productivity Center.• IBM Tivoli Storage Productivity Center for Fabric Supported Products. Click this item to view the Support for Tivoli Storage Productivity Center for Fabric Web page. To view the devices supported by Fabric Manager, scroll to the Integration/Operation section and click the link that matches your version of Tivoli Storage Productivity Center.• IBM Tivoli Storage Productivity Center for Data Supported Products. Click this item to view the Support for Tivoli Storage Productivity Center for Data Web page. To view the devices supported by Disk Manager, scroll to the Integration/Operation section and click the link that matches your version of Tivoli Storage Productivity Center. <p>See "Web browser support" on page 45 for information about Web browser support for launching the supported products data Web pages.</p>

Menu Option	Description
Tivoli Storage Productivity Center Demos	Select this menu item to launch a Web page that contains Flash demos of Tivoli Storage Productivity Center. These demos provide an introduction to the product, the business and storage administrator challenges that are met by the product, and detailed information about the features and functions that the product provides. See “Web browser support” for information about Web browser support for launching the product demos Web page.
Launch IBM Tivoli Storage Productivity Center Information Center	Select this menu item to display product information including: <ul style="list-style-type: none"> • version information • copyright information • trademark information

Web browser support:

IBM Tivoli Storage Productivity Center starts a web browser when you access web pages from items in the Help menu, launch another application using its launch-in-context feature, or start IBM Tivoli Storage Productivity Center for Replication from its user interface.

Note: The Tivoli Storage Productivity Center software is bundled with its own IBM Java JRE. When you use the Java Web Start method of remotely accessing the Tivoli Storage Productivity Center GUI, you must use the IBM Java.

You can use the following web browsers with Tivoli Storage Productivity Center:

AIX

- Mozilla Firefox 2.0
- Mozilla Firefox 3.0

Linux

- Mozilla Firefox 2.0
- Mozilla Firefox 3.0

Windows

- Internet Explorer 7
- Internet Explorer 8 (in Compatibility View)

In Internet Explorer, the Computability View option allows you to display a web page that was designed for a version older than version 8 to be displayed correctly in Internet Explorer 8. To display your web page in Internet Explorer 8.0, click **Tools > Compatibility View Settings**.

- Mozilla Firefox 2.0
- Mozilla Firefox 3.0

Note:

- To get Firefox 2 or 3 for Linux or Windows, go to <ftp://ftp.mozilla.org/pub/mozilla.org/firefox/releases/>.
- To get Firefox version 2 on AIX, you can use a previous version of IBM Firefox Browser for AIX CD (if you have it). Or you can download the source code for version 2 or 3 from <ftp://ftp.mozilla.org/pub/mozilla.org/firefox/releases/>.

- If you do not have a web browser configured for use with Tivoli Storage Productivity Center when you access a web page from its user interface, the Configure Element Launcher window is displayed. Use this window to configure a web browser for use with the product.
- For information about the web browsers supported for the DS8000, see the DS8000 Information Center at <http://publib.boulder.ibm.com/infocenter/dsichelp/ds8000ic/index.jsp>. Search for **Internet browser support**.
- For information about launching Java Web Start using Internet Explorer 8, go to the information center. Search for *Starting the Tivoli Storage Productivity Center GUI using a web browser*.

The pop-up blocker must be turned off in Windows Internet Explorer for any application to start the graphical user interface (GUI) for the following storage systems:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

For example, if you are in Tivoli Storage Productivity Center and you want to start the Element Manager of the SAN Volume Controller GUI, you must configure Internet Explorer with the following settings.

1. Open Internet Explorer.
2. From the Internet Explorer toolbar, click **Tools > Pop-up Blocker > Turn Off Pop-up Blocker**.

If you receive a message that content has been blocked because it was not signed by a valid security certificate, click the Information Bar at the top of the window and select **Show blocked content**.

Web browser support for Tivoli Storage Productivity Center for Replication

When you start Tivoli Storage Productivity Center for Replication from within the Tivoli Storage Productivity Center, its user interface appears in a web browser. Depending on the configuration settings for the web browser on your computer, Tivoli Storage Productivity Center for Replication appears within a tab of an existing browser session or it appears within a new browser session.

In Firefox, you can configure how Tivoli Storage Productivity Center for Replication appears when launched from Tivoli Storage Productivity Center. To configure Firefox, complete the following steps:

1. Start Firefox.
2. Select **Tools > Options** from the menu bar.
3. Select **Tabs**.
4. Select **a new window** to indicate that you want start Tivoli Storage Productivity Center for Replication in a new web browser session. Select **a new tab** to indicate that you want to start Tivoli Storage Productivity Center for Replication in a new tab of an existing web browser session.

Note: For Internet Explorer, Tivoli Storage Productivity Center for Replication is always started in a new web browser session.

Note: If you start Tivoli Storage Productivity Center for Replication from within Tivoli Storage Productivity Center, you might be logged out of the Tivoli Storage Productivity Center for Replication user interface unexpectedly. This action occurs

when you use Tivoli Storage Productivity Center to start a session of Tivoli Storage Productivity Center for Replication in an Internet Explorer web browser and then open and close a wizard or secondary window in that Tivoli Storage Productivity Center for Replication session.

To work around this problem, use one of these solutions:

- Start Tivoli Storage Productivity Center for Replication from Tivoli Storage Productivity Center using Firefox
- Start Tivoli Storage Productivity Center for Replication from Tivoli Integrated Portal
- Start Tivoli Storage Productivity Center for Replication by entering its web address directly in a Firefox or Internet Explorer web browser.

Web browser support for Help menu items

Tivoli Storage Productivity Center provides a number of items in its **Help** menu that you can click to learn more about how to use the product. When you click one of these items, Tivoli Storage Productivity Center launches a web browser and displays the appropriate web page. You must have an internet connection on the system where you are running the Tivoli Storage Productivity Center user interface. You must also have a default browser configured to successfully access these web pages.

Toolbar



The toolbar at the top of the Tivoli Storage Productivity Center main window provides one-click access to key commands.









TVT Instructions:btmtoolb.jpg changed for V4.2:

Start and log in to Tivoli Storage Productivity Center. The toolbar is displayed on the main window.

Click the following buttons on the toolbar to access key commands in the product:

Button	Description
Element Management	Click this button to open the Element Manager page in the Configuration Utility. Use the Element Manager page to configure and start the element managers that are visible to Tivoli Storage Productivity Center. Element managers are programs that enable you to configure and maintain your storage devices. An example of an element manager is DS8000 Storage Manager, which is the default element manager that is provided with IBM System Storage DS8000.
Back 	Click this button to move backward through the active panels in the content pane.
Forward 	Click this button to move forward through the active panels in the content pane.

Button	Description
Save 	Click this button to save the current object that you are editing. The object changes are saved to the database repository.
Print 	Click this button to print the current data or report chart that is displayed in the content pane.
Cancel edit session 	Click this button to discard the contents of the content pane. If content pane is in an edit session and changes have been made, a prompt notifying you that a change has been made is displayed.
Configure devices 	<p>Access the Configure Devices wizard. Use the wizard to configure the following storage devices for monitoring by Tivoli Storage Productivity Center.</p> <ul style="list-style-type: none"> • Storage systems • Fabrics and switches • Computers • Tape libraries <p>The wizard guides you through the process of adding a storage device as a data source, running a discovery, applying alert conditions, and including the device in data collection schedules.</p>
Manage jobs, schedules, and runs 	Access the Job Management panel. Use this panel to view and manage the schedules, runs, and jobs related to the storage devices that are monitored by Tivoli Storage Productivity Center.
Stop current server request 	<p>Click this button to stop the current server request. This button becomes active when a request is made to the server and the user interface is waiting for a response.</p> <p>During the server request wait time, the cursor arrow is displayed as an hourglass. If you decide the current request is taking too long or you do not want to wait any longer, click this button to cancel that request and regain control of the GUI.</p> <p>Note: If you click this button, the current request continues to run, and the resources that are used by it on the server are not released until it completes. Frequent use of this button during a session might exhaust the resources on the server and cause the user interface to stop responding.</p>

Schedules, runs, and jobs

Use schedules, run, and jobs to manage and monitor the storage entities in your environment. These objects form the framework of job management within Tivoli Storage Productivity Center. Use the Job Management panel to view and manage the schedules, runs, and jobs related to the storage entities that are monitored by the product.

In previous versions of Tivoli Storage Productivity Center, a schedule and its related runs and jobs were displayed as nodes in different parts of the navigation tree. In v4.2 and later, all schedules, runs, and jobs are displayed in the Job Management panel. Use this panel to perform the following tasks:

- View information about all schedules, runs, and jobs in one central location.

- Identify all the schedules, runs, and jobs that are encountering problems when executing.
- Ensure that a storage entity is fully monitored by Tivoli Storage Productivity Center.
- Edit, rename, delete, and run schedules.

Schedules

A schedule is a definition of a monitoring task that collects data about storage devices. Monitoring tasks include scans, probes, pings, quotas, discovery, and performance monitors. For example, define a probe schedule to collect data about storage subsystems; define a scan schedule to collect file-level information about hosts.

Schedule definitions are displayed in the navigation tree under their corresponding type. For example, probe schedules are displayed under the **IBM Tivoli Storage Productivity Center > Monitoring > Probes** node; scans are displayed under the **Data Manager > Monitoring > Scans** node. You can click a schedule in the navigation tree to view or edit its definition at any time.

Runs A run is a single invocation of a schedule. For example, if a probe schedule is set to collect data once per day, there is one run of that schedule per day. Runs for a schedule are displayed on the Job Management panel.

Jobs A job is a unit of work within a run for a schedule. One run might have multiple jobs. For example, a subsystem probe schedule has one job created for each subsystem for which it collects data. Jobs for a run are displayed on the Job Management panel.

Tivoli Storage Productivity Center for Replication

This topic provides a brief overview of IBM Tivoli Storage Productivity Center for Replication.

Tivoli Storage Productivity Center includes Tivoli Storage Productivity Center for Replication. The Tivoli Storage Productivity Center for Replication FlashCopy and Snapshot replication features are included with all editions of Tivoli Storage Productivity Center. Other Tivoli Storage Productivity Center for Replication replication features require one of the following licenses:

- Tivoli Storage Productivity Center for Replication Two Site Business Continuity (BC): Metro Mirror and Global Mirror replication.
- Tivoli Storage Productivity Center for Replication Three Site BC: Metro Mirror, Global Mirror, and Metro Global Mirror to secondary and tertiary sites.

Tivoli Storage Productivity Center for Replication provides you with continuous availability and disaster recovery solutions with both point-in-time replication, and continuous replication. Other options include the ability to practice what you would do in the event of a disaster, and management servers that enable you to manage your disaster recovery solution if you experience a site outage.

Tivoli Storage Productivity Center for Replication Two Site Business Continuity

You can use Tivoli Storage Productivity Center for Replication Two Site Business Continuity to obtain continuous availability and disaster recovery solutions by using point-in-time replication, which includes FlashCopy and Snapshot, and continuous replication, which includes Metro Mirror and Global Mirror.

You can set up Metro Mirror and Global Mirror sessions to replicate data in both the forward and reverse directions.

Use Tivoli Storage Productivity Center for Replication Two Site Business Continuity to create and manage the following session types:

- FlashCopy:
 - FlashCopy replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server Model 800
 - FlashCopy replication for SAN Volume Controller
 - FlashCopy replication for Storwize V7000
 - FlashCopy replication for Storwize V7000 Unified
- Snapshot:
 - Snapshot replication for the XIV system
- Global Mirror:
 - Global Mirror Single Direction replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
 - Global Mirror Single Direction replication for SAN Volume Controller
 - Global Mirror Single Direction replication for Storwize V7000
 - Global Mirror Single Direction replication for Storwize V7000 Unified
 - Global Mirror Failover/Failback replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
 - Global Mirror Failover/Failback replication for SAN Volume Controller
 - Global Mirror Failover/Failback replication for Storwize V7000
 - Global Mirror Failover/Failback replication for Storwize V7000 Unified
 - Global Mirror Failover/Failback replication for XIV system
 - Global Mirror Failover/Failback with Practice replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
 - Global Mirror Failover/Failback with Practice replication for SAN Volume Controller
 - Global Mirror Failover/Failback with Practice replication for Storwize V7000
 - Global Mirror Failover/Failback with Practice replication for Storwize V7000 Unified
 - Global Mirror Either Direction with Two-Site Practice replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
- Metro Mirror:
 - Metro Mirror or Global Copy replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
 - Metro Mirror for Single Direction replication for SAN Volume Controller
 - Metro Mirror for Single Direction replication for Storwize V7000
 - Metro Mirror for Single Direction replication for Storwize V7000 Unified
 - Metro Mirror for Failover/Failback replication for SAN Volume Controller
 - Metro Mirror for Failover/Failback replication for Storwize V7000
 - Metro Mirror for Failover/Failback replication for Storwize V7000 Unified
 - Metro Mirror Failover/Failback replication for XIV system
 - Metro Mirror for Failover/Failback with Practice replication for SAN Volume Controller

- Metro Mirror for Failover/Failback with Practice replication for Storwize V7000
- Metro Mirror for Failover/Failback with Practice replication for Storwize V7000 Unified
- Metro Mirror Failover/Failback or Global Copy replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
- Metro Mirror Failover/Failback replication for HyperSwap® for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
- Metro Mirror Failover/Failback replication for Open HyperSwap for the System Storage DS8000

Tivoli Storage Productivity Center for Replication Three Site Business Continuity

You can use Tivoli Storage Productivity Center for Replication Three Site Business Continuity to obtain continuous availability and disaster recovery solutions by using point-in-time replication, which includes FlashCopy and Snapshot, and continuous replication, which includes Metro Mirror, Global Mirror, and Metro Global Mirror to secondary and tertiary sites.

You can set up Metro Mirror and Global Mirror sessions to replicate data in both the forward and reverse directions. With Tivoli Storage Productivity Center for Replication Three Site Business Continuity, you can also use Metro Global Mirror (with failover and failback to switch production sites between each of the three sites and return to the original configuration).

Important: Tivoli Storage Productivity Center for Replication Three Site Business Continuity requires that Tivoli Storage Productivity Center for Replication Two Site Business Continuity is installed and a separate license is required for both products.

Use Tivoli Storage Productivity Center for Replication Three Site Business Continuity to create and manage the following session types:

- FlashCopy:
 - FlashCopy replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server Model 800
 - FlashCopy replication for SAN Volume Controller
 - FlashCopy replication for Storwize V7000
 - FlashCopy replication for Storwize V7000 Unified
- Snapshot:
 - Snapshot replication for the XIV system
- Global Mirror:
 - Global Mirror Single Direction replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
 - Global Mirror Single Direction replication for SAN Volume Controller
 - Global Mirror Single Direction replication for Storwize V7000
 - Global Mirror Single Direction replication for Storwize V7000 Unified
 - Global Mirror Failover/Failback replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
 - Global Mirror Failover/Failback replication for SAN Volume Controller

- Global Mirror Failover/Failback replication for Storwize V7000
- Global Mirror Failover/Failback replication for Storwize V7000 Unified
- Global Mirror Failover/Failback replication for XIV system
- Global Mirror Failover/Failback with Practice replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
- Global Mirror Failover/Failback with Practice replication for SAN Volume Controller
- Global Mirror Failover/Failback with Practice replication for Storwize V7000
- Global Mirror Failover/Failback with Practice replication for Storwize V7000 Unified
- Global Mirror Either Direction with Two-Site Practice replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
- Metro Mirror:
 - Metro Mirror or Global Copy replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
 - Metro Mirror for Single Direction replication for SAN Volume Controller
 - Metro Mirror for Single Direction replication for Storwize V7000
 - Metro Mirror for Single Direction replication for Storwize V7000 Unified
 - Metro Mirror for Failover/Failback replication for SAN Volume Controller
 - Metro Mirror for Failover/Failback replication for Storwize V7000
 - Metro Mirror for Failover/Failback replication for Storwize V7000 Unified
 - Metro Mirror Failover/Failback replication for XIV system
 - Metro Mirror for Failover/Failback with Practice replication for SAN Volume Controller
 - Metro Mirror for Failover/Failback with Practice replication for Storwize V7000
 - Metro Mirror for Failover/Failback with Practice replication for Storwize V7000 Unified
 - Metro Mirror Failover/Failback or Global Copy replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
 - Metro Mirror Failover/Failback with HyperSwap replication for System Storage DS8000, System Storage DS6000, and TotalStorage Enterprise Storage Server 800
 - Metro Mirror Failover/Failback with Open HyperSwap replication for System Storage DS8000
- Metro Global Mirror:
 - Metro Global Mirror replication for System Storage DS8000 (with failover and failback)
 - Metro Global Mirror replication for System Storage DS8000 and TotalStorage Enterprise Storage Server 800 (with failover and failback)
 - Metro Global Mirror with Practice replication for System Storage DS8000 and TotalStorage Enterprise Storage Server 800 (with failover and failback)
 - Metro Global Mirror with HyperSwap replication for System Storage DS8000 and TotalStorage Enterprise Storage Server 800 (TotalStorage Enterprise Storage Server 800 can be only in the H1 role.)

Tivoli Storage Productivity Center for Replication integration with Tivoli System Automation Application Manager

IBM Tivoli System Automation Application Manager includes a Distributed Disaster Recovery (DDR) feature that uses IBM Tivoli Storage Productivity Center for Replication to provide a data replication and disaster recovery solution for distributed systems. The DDR feature is available in Tivoli System Automation Application Manager version 3.2 or later and requires Tivoli Storage Productivity Center for Replication version 3.4.1.8 or version 4.1.1.1 or later.

Tivoli System Automation Application Manager automates the availability of information technology resources. The DDR feature and Tivoli Storage Productivity Center for Replication provide a full end-to-end solution for site switching on distributed systems at the application level.

For more information about the integration of Tivoli Storage Productivity Center for Replication with Tivoli System Automation Application Manager, including the steps required to perform the integration, see the *Tivoli System Automation Application Manager Administrator's and User's Guide*. The administrator's guide is provided in the Tivoli System Automation Application Manager information centers by version at <http://www.ibm.com/developerworks/wikis/display/tivolidoccentral/Tivoli+System+Automation+Application+Manager>.

Tivoli Integrated Portal

The IBM Tivoli Storage Productivity Center installation program includes IBM Tivoli Integrated Portal. Tivoli Integrated Portal is a standards-based architecture for Web administration. Tivoli Integrated Portal enables developers to build administrative interfaces for IBM and independent software products as individual plug-ins to a common console network. The installation of Tivoli Integrated Portal is required to enable single sign-on for Tivoli Storage Productivity Center.

Tivoli Integrated Portal and single sign-on

Single sign-on is an authentication process that you can use to enter one user ID and password to access multiple applications. For example, you can access Tivoli Integrated Portal and then access Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication from Tivoli Integrated Portal using a single user ID and password. Single sign-on integrates with the launch and launch in context features to enable you to move smoothly from one application to another application.

Related concepts:

“Single sign-on” on page 58

Single sign-on is an authentication process that enables you to enter one user ID and password to access multiple applications. For example, you can access IBM Tivoli Integrated Portal and then access IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication from Tivoli Integrated Portal using a single user ID and password. Single sign-on integrates with the launch and launch in context features to enable you to move smoothly from one application to another application.

“Launch and launch in context” on page 60

The launch and launch in context features are used to start IBM Tivoli Storage Productivity Center and to start other applications from Tivoli Storage Productivity Center. The terms inbound and outbound are used to differentiate starting Tivoli Storage Productivity Center (inbound) and starting another application from Tivoli Storage Productivity Center (outbound).

Key concepts

This section contains a technical overview that will help you understand how IBM Tivoli Storage Productivity Center works. An understanding of the concepts in this section will help you use Tivoli Storage Productivity Center effectively.

Starting and stopping the Tivoli Storage Productivity Center GUI

This topic describes how to start and stop the Tivoli Storage Productivity Center graphical user interface (GUI).

Starting the Tivoli Storage Productivity Center GUI

Learn how to start the Tivoli Storage Productivity Center graphical user interface (GUI) from one of the following locations: the Microsoft Windows Start menu, IBM Tivoli Integrated Portal, the Productivity Center icon on your desktop (Windows), or the command line (AIX, UNIX, or Linux).

Starting the Tivoli Storage Productivity Center GUI from Tivoli Integrated Portal:

You can start the Tivoli Storage Productivity Center graphical user interface (GUI) from Tivoli Integrated Portal. With the single sign-on feature, you can log on to the Tivoli Storage Productivity Center GUI without having to explicitly enter a user name and password.

To start the Tivoli Integrated Portal and Tivoli Storage Productivity Center, complete the following steps:

1. Ensure that you are using these browsers:

AIX Firefox 2.0 or 3.0

Linux and UNIX
Firefox 2.0 or 3.0

Windows
Internet Explorer 7 or 8
Firefox 2.0 or 3.0

2. Ensure that a fully qualified domain name is defined for the Tivoli Storage Productivity Center server. To verify the Tivoli Storage Productivity Center

server name and modify the name if required, go to the information center. Search for **Data server and Device server are inaccessible from Tivoli Integrated Portal**. Contact your Tivoli Storage Productivity Center administrator before you modify the file.

3. If you are using Internet Explorer 8 the first time, click **Start Tivoli Storage Productivity Center**.
4. Click the message and select **Download file**. You can then start Java Web Start.
5. To start Tivoli Storage Productivity Center from Tivoli Integrated Portal, complete the following steps:

- a. Start an Internet Explorer or Firefox Web browser and type the following information in the address bar:

`http://hostname:port`

where *hostname* defines the server that is running Tivoli Integrated Portal such as the server name or IP address and *port* defines the port number for Tivoli Integrated Portal. If the default port was accepted during the installation of Tivoli Integrated Portal, the port number is 16310. Contact your Tivoli Storage Productivity Center administrator if you need to verify the host name and port number

- b. On the Tivoli Integrated Portal logon page, log on using the appropriate user ID and password. Your user ID must have administrator permissions.
- c. In the Tivoli Integrated Portal navigation tree, click **Tivoli Storage Productivity Center**
- d. On the Tivoli Storage Productivity Center portlets page, click **Start Storage Productivity Center**, and one of the following occurs:
 - If single sign-on is successful, Tivoli Storage Productivity Center starts without displaying the logon window.
 - If single sign-on is not successful, an error message and the Tivoli Storage Productivity Center logon window are displayed.
 - If the Tivoli Storage Productivity Center Data server or Device server is not accessible, an error message is displayed. The status of both the Data server and the Device server also display.
 - If the status of the Device server and Data server is inaccessible and you started Tivoli Integrated Portal from a remote computer, it is possible that a fully qualified domain name was not defined for the Tivoli Storage Productivity Center server during installation. To check the Tivoli Storage Productivity Center server name and modify the name if required, see the Tivoli Storage Productivity Center troubleshooting documentation.
 - If you are using a Lightweight Directory Access Protocol (LDAP) compliant directory for Tivoli Storage Productivity Center user authentication and the directory is not available, an error message is displayed.

Once you have logged on to Tivoli Integrated Portal, a Lightweight Third-Party Authentication (LTPA) token is created. This token is passed to other applications that you start from Tivoli Integrated Portal for single sign-on authentication purposes.

During the period between when you log on to Tivoli Integrated Portal and when you start another application such as Tivoli Storage Productivity Center from Tivoli Integrated Portal, the following conditions might occur:

- The user password that was used to log on to Tivoli Integrated Portal is changed in the user repository.

- The user ID that was used to access Tivoli Integrated Portal is changed in the repository or removed from the user repository.
- The user repository is not accessible.

Under the first condition, the original user credentials that were used to access Tivoli Integrated Portal are used to access other applications until the timeout period for the LTPA token that is used for single sign-on expires. When the LTPA token expires, you are prompted to re-enter your user ID and password when you attempt to start another application using single sign-on.

Under the second and third conditions, the single sign-on feature does not work. You are always prompted to re-enter your user ID and password when you attempt to start another application.

Starting the Tivoli Storage Productivity Center GUI as a stand-alone application:

This section describes how to start the IBM Tivoli Storage Productivity Center from the Microsoft Windows Start menu, or the Productivity Center icon on your desktop (Windows), or from the command line (AIX, UNIX, or Linux).

To start the Tivoli Storage Productivity Center graphical user interface (GUI) on Windows, click **Start > Programs > IBM Tivoli Storage Productivity Center > Productivity Center**. You can also double-click the IBM Tivoli Storage Productivity Center icon if it is installed on your desktop.

To start the Tivoli Storage Productivity Center GUI on UNIX or Linux, type the following path and command at the command line:

```
/opt/IBM/TPC/gui/TPCD.sh
```

To start the Tivoli Storage Productivity Center GUI on AIX, type the following path and command at the command line:

```
/opt/IBM/TPC/gui/TPCD.sh
```

Logging on to Tivoli Storage Productivity Center:

Use the logon window to specify the user ID and password that you want to use to log on to a IBM Tivoli Storage Productivity Center server.

Define logon information for Tivoli Storage Productivity Center in the following fields:

User ID

Enter the user ID that you want to use to log on to Tivoli Storage Productivity Center. The roles that are assigned to that user ID determine what nodes in the navigation tree you can see and act upon.

Password

Enter the password that is associated with the **User ID**.

Server Enter the IP address or Domain Name System (DNS) name of the computer on which the Tivoli Storage Productivity Center server is installed.

When you enter an IP address, you can use either an IPv4 or IPv6 address format depending on the protocol enabled on the computer. For example, enter an IPv6 address for an IPv6-only computer. Enter an IPv4 or IPv6 address for a computer with dual stacks (both IPv4 and IPv6 enabled).

- To enter an IPv4 address, use the dotted decimal format: *nnn.nnn.nnn.nnn*. For example: 127.0.0.1.
- When entering an IPv6 address you must include brackets [] to separate those addresses from their port numbers when you enter a value in this field. For example: [2001:0DB8::1234:0000:0000:5678:ABCD]:9550. Use one of the following methods to enter an IPv6 address:
 - The preferred IPv6 address representation is written as eight groups of four hexadecimal digits: *xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx*, where each *x* is a hexadecimal digit representing 4 bits. For example: 2001:0DB8:0000:1234:0000:0000:5678:ABCD
 - You can specify an IPv6 address using a shortened format that omits leading zeros: (2001:DB8:0:1234:0:0:5678:ABCD)
 - You can use double colons in place of a series of zeros: (2001:DB8:0000:1234::5678:ABCD)

Setting up a user role and collecting system statistics

If you are logging on to the Tivoli Storage Productivity Center graphical user interface (GUI) for the first time after installing the product, note the following information:

- If you did not set up a Tivoli Storage Productivity Center user role, you might not be able to log on. To set up a Tivoli Storage Productivity Center user role, log on as the superuser and create the roles as described in the IBM Tivoli Storage Productivity Center Information Center.
- After you log on, you will see a panel that prompts you to collect system statistics. Click **Yes**.

Single sign-on

Tivoli Storage Productivity Center can use the single sign-on feature, which enables you to start Tivoli Storage Productivity Center from either the IBM Tivoli Integrated Portal instance that is installed with Tivoli Storage Productivity Center or from an existing Tivoli Integrated Portal V1.1.1.x instance. If you have a single sign-on environment configured and Tivoli Integrated Portal user authentication is successful, the Tivoli Storage Productivity Center logon window does not open when you start the Tivoli Storage Productivity Center GUI from Tivoli Integrated Portal.

User credential considerations for single sign-on

When you log on to either Tivoli Integrated Portal or Tivoli Storage Productivity Center, a Lightweight Third-Party Authentication (LTPA) token is created and used for single sign-on authentication when you start other applications from within the Tivoli Storage Productivity Center GUI.

During the period between when you log on to Tivoli Storage Productivity Center and when you start other applications, such as element managers, from the Tivoli Storage Productivity Center GUI, the following conditions might occur:

The user password that was used to log on to Tivoli Integrated Portal or Tivoli Storage Productivity Center is changed in the user repository.	Under this condition, the original user credentials that were used to log into Tivoli Integrated Portal or Tivoli Storage Productivity Center are used to access other applications until the timeout period for the LTPA token that is used for single sign-on expires. When the LTPA token expires, you are prompted to reenter your user ID and password when you attempt to start another application using single sign-on.
The user ID that was used to access Tivoli Integrated Portal is changed in the repository or removed from the user repository.	Under the second and third conditions, the single sign-on feature does not work. You are always prompted to re-enter your user ID and password when you attempt to start another application through single sign-on.
The user repository is not accessible.	

Stopping the Tivoli Storage Productivity Center GUI

This section describes how to stop the IBM Tivoli Storage Productivity Center graphical user interface (GUI).

To stop the Tivoli Storage Productivity Center user interface, click **File > Exit**.

Single sign-on

Single sign-on is an authentication process that enables you to enter one user ID and password to access multiple applications. For example, you can access IBM Tivoli Integrated Portal and then access IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication from Tivoli Integrated Portal using a single user ID and password. Single sign-on integrates with the launch and launch in context features to enable you to move smoothly from one application to another application.

For the current version of Tivoli Storage Productivity Center, single sign-on is available for the following applications:

- Tivoli Storage Productivity Center for Replication
- IBM System Storage DS8000
- IBM System Storage SAN Volume Controller
- IBM Storwize V7000

For Tivoli Storage Productivity Center for Replication Storage Manager, single sign-on is enabled by default.

For DS8000 Storage Manager, SAN Volume Controller, and Storwize V7000, single sign-on must be enabled.

To enable single sign-on, see the following topics in the Tivoli Storage Productivity Center information center:

- Configuring Tivoli Storage Productivity Center for DS8000 LDAP authentication
- Configuring Tivoli Storage Productivity Center for SAN Volume Controller and Storwize V7000 single sign-on

How single sign-on works

A single sign-on environment requires a centralized authentication repository that is accessed by all applications within the environment. The user's ID and other authentication information are passed between applications using Lightweight Third-Party Authentication (LTPA) tokens. LTPA is the security technology that is used in the IBM WebSphere Application Server for passing the user authentication information between applications. To use the LTPA tokens, each application must have the same set of public keys to encrypt the user's information. The authenticating service uses the corresponding private keys to decrypt the user's information for authentication. As an additional security mechanism, LTPA tokens expire after a determined amount of time. By default, the tokens expire after 24 hours. You can change the LTPA token expiration time using the Tivoli Integrated Portal user interface.

Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication use their respective IBM WebSphere Application Server instances to authenticate LTPA tokens. However, other applications, such as the IBM System Storage DS8000 element manager and other element managers that do not run within an IBM WebSphere Application Server instance use the authentication service that is provided with Tivoli Integrated Portal. The authentication service client is typically embedded in these other applications and the client communicates with the authentication service server in Tivoli Integrated Portal for all authentication requests.

The single sign-on feature is not supported by the Tivoli Storage Productivity Center command-line interface (CLI).

Selecting the user authentication method to use with single sign-on

During the installation of Tivoli Storage Productivity Center, you can specify whether to use LDAP or the operating system as the authentication and authorization repository (see the following description of each authentication type). If you do not specify LDAP, then Tivoli Storage Productivity Center uses the operating system (OS) users and groups on the server where Tivoli Storage Productivity Center is installed for authentication and authorization.

OS Authentication

This method authenticates the user against the users defined for the local operating system (OS).

LDAP/Active Directory

This method authenticates the user against a Lightweight Directory Access Protocol (LDAP) or Microsoft Active Directory repository.

If OS authentication is selected, the use of the single sign-on feature is limited. OS authentication does not support single sign-on for element managers, even when the element manager is installed on the same computer as Tivoli Storage Productivity Center. LDAP or Microsoft Active Directory authentication supports single sign-on for element managers regardless of where they are installed.

You can change the user authentication method using the Tivoli Integrated Portal graphical user interface (GUI).

During installation, you can specify that Tivoli Storage Productivity Center use an existing Tivoli Integrated Portal instance rather than the instance that is installed

with Tivoli Storage Productivity Center. In this case, Tivoli Storage Productivity Center uses the authentication repository that is used by Tivoli Integrated Portal (LDAP or the operating system). However, if the existing Tivoli Integrated Portal instance uses a file-based authentication repository, you cannot use the existing Tivoli Integrated Portal instance. You must install a new Tivoli Integrated Portal instance when you install Tivoli Storage Productivity Center.

Related concepts:

“Changing the user authentication method” on page 225

The IBM Tivoli Storage Productivity Center installation program enables you to select a user authentication method that is used by Tivoli Storage Productivity Center, IBM Tivoli Storage Productivity Center for Replication, and IBM Tivoli Integrated Portal. You can choose to authenticate users against the users defined for the local operating system or to authenticate users against the users defined in a Lightweight Directory Access Protocol (LDAP) or Microsoft Active Directory repository. You can change the user authentication method after installation using Tivoli Integrated Portal.

Related tasks:

“Changing the LTPA token expiration for single sign-on” on page 242

For single sign-on, the user's authentication information is passed between applications using Lightweight Third-Party Authentication (LTPA) tokens. By default, the tokens expire after 24 hours. You can change the LTPA token expiration time using the Tivoli Integrated Portal user interface.

Launch and launch in context

The launch and launch in context features are used to start IBM Tivoli Storage Productivity Center and to start other applications from Tivoli Storage Productivity Center. The terms inbound and outbound are used to differentiate starting Tivoli Storage Productivity Center (inbound) and starting another application from Tivoli Storage Productivity Center (outbound).

Inbound launch in context refers to starting Tivoli Storage Productivity Center at a specific point in the graphical user interface (GUI). You can use inbound launch in context in either of the following ways:

- Configure an application such as Tivoli Enterprise Portal to start Tivoli Storage Productivity Center at functionally related locations in the Tivoli Storage Productivity Center GUI.
- Start Tivoli Storage Productivity Center from Java Web Start, the command line, or the Windows Start menu using parameters to specify the starting location in the Tivoli Storage Productivity Center GUI.

Outbound launch in context refers to starting other applications, such as administrative applications for storage systems and switches, from functionally related points in the Tivoli Storage Productivity Center GUI.

Launch refers to starting an application at the starting point for that application. Launch primarily refers to starting applications from the Tivoli Storage Productivity Center GUI. Starting applications from Tivoli Storage Productivity Center is referred to as an outbound launch.

Related concepts:

“Inbound launch in context feature”

With the inbound launch in context feature you can start IBM Tivoli Storage Productivity Center at a specified point in the graphical user interface (GUI).

“**Outbound launch and launch in context features**” on page 67

With the outbound launch and launch in context features, you can start the administrative applications for devices such as storage systems and switches from various points in the graphical user interface (GUI).

Inbound launch in context feature

With the inbound launch in context feature you can start IBM Tivoli Storage Productivity Center at a specified point in the graphical user interface (GUI).

Prerequisite: You must have View and Modify permissions to use the launch in context feature.

You can use the launch in context feature to perform the following tasks in Tivoli Storage Productivity Center:

- Create a volume or VDisk
- Open the storage area network SAN Planner
- Create a capacity report
- Create a wasted space report

You can run the launch in context feature in the following ways:

- **Use Java Web Start** from a Web browser. See “Launch with Java Web Start.”
- **Use the command line on Windows, AIX, Linux, or UNIX** to enter the command and arguments on the command line. See “Launch from the command line” on page 63.
- **Use the Windows Start menu.** See “Launch from Start menu” on page 64.
- **Use the Tivoli Enterprise Portal GUI** to define, edit, and launch commands. See “Launch from Tivoli Enterprise Portal” on page 65.

The launch in context feature has a number of parameters. See “Launch in context parameters” on page 1151, “Launch in context task parameters” on page 1153, and “Available windows” on page 1154.

Launch with Java Web Start:

Using Java Web Start, you can use the launch in context feature from a Uniform Resource Locator (URL).

The launch in context feature requires Java Web Start 6 on the computer that is running the Web browser. If the correct release of Java Web Start is not present, you receive an error message when you enter the URL. The message states that you must download the Java Runtime Environment, which includes Java Web Start. Follow the link provided in the message to install the required software.

If you are using Internet Protocol Version 6 (IPv6) and are using Java Web Start to start IBM Tivoli Storage Productivity Center on a Windows system, use a system host name such as **tpc332.storage.newyork.xyz.com:9549** in the Web browser address field rather than an IPv6 address such as **[2001:DB8:0000:1234:0000:0000:5678:ABCD]:9549**.

To use the launch in context feature to perform a task or display a window on a remote server from your browser, complete the following steps:

1. Open your browser and type the launch in context command.

URLs cannot include spaces. In the command, a parameter and its value are separated by an equal sign (=). Parameter=value pairs are separated by an ampersand (&). If any of the parameters are spelled incorrectly, an error message indicates the incorrect parameter.

The format of a command to create a volume is:

```
https://device_server:device_port/ITSRM/jsp/jnlp/tpcgui.jsp?consoleClass=com.ibm.usmi.console.lic.tpc.LICConsole&-server=server:port&-function=disk.create_volume&-subsystem_id=ESS-2105-17844-IBM&-name=myvol&-number=2&-size=10&-unit=MB
```

The format of a command to view a window on a remote server is:

```
http://device_server:device_port/ITSRM/jsp/jnlp/tpcgui.jsp?consoleClass=com.ibm.usmi.console.lic.tpc.LICConsole&-server=server:port&-function=window_name
```

Tip: The Data server port value is 9549. The port value for the Device server is 9551 for https and 9550 for http. If the wrong the Device server port value is used, the return page cannot be interpreted by the browser. If the wrong Data server port value is used, the following error message is returned.

Unable to connect to the server

2. The first time a window or task is started, you receive a security warning informing you that the downloaded jar files have been digitally signed using a certificate from a trusted company. To prevent this message from being displayed in the future, select **Always trust content from this publisher** and click **Run**. If the Tivoli Storage Productivity Center graphical user interface (GUI) has not been downloaded on the local computer using Java Web Start, or the local computer does not have the latest version of the Tivoli Storage Productivity Center GUI, the GUI is downloaded.
3. If you included the user ID and password in the command, the sign-in dialog is bypassed, otherwise log on to the Tivoli Storage Productivity Center GUI. The Tivoli Storage Productivity Center GUI navigates to the tree node corresponding to the window name that you specified. You are connected using the user name and password that you entered. Depending on the value you specified for the -function parameter, the following actions might occur:
 - The specified window opens
 - For disk.create_volume, depending on the parameters, the Create Volume wizard or Create VDisk wizard opens.
 - For disk.san_planner, the SAN Planner panel opens with the given SAN plan ID.
 - For data.filesystem_capacity_by_computer, the Capacity Report opens.
 - For TPC.reports_data_wasted_space, the Wasted Space Report opens.

Restrictions:

- When you open a plan in the SAN Planner using launch in context, you must enter the exact name of the plan. The value for the **-plan_id** parameter is case-sensitive and must match the name of the plan when it was created. For example, if you created a plan named **UserPlan1**, you must specify **-plan_id UserPlan1** in the launch in context command.
- When you launch the Filesystem Capacity > By Computer report using launch in context, you must specify the ID for the computer upon which you want to report. For example:

```
http://tpcserver14:9550/ITSRM/jsp/jnlp/
tpcgui.jsp?consoleClass=com.ibm.usmi.console.lic.tpc.LICConsole&-
server=tpcserver14.storage.xyz.com:9549&-
function=data.filesystem_capacity_by_computer&-computer_id=tpcserver14
```

The value you enter for the **computer_id** parameter might not match the ID for that computer that is stored in the Tivoli Storage Productivity Center database repository. You receive a warning message if the ID for the computer that you enter in the URL does not match the ID that is stored in the data repository. To ensure that this report is displayed properly when you run from the launch in context facility, enter an ID for the computer that matches the ID that is displayed for it through the Topology Viewer or appropriate report in the product user interface.

Use the information in the links below to enter parameters.

“Launch in context parameters” on page 1151

You can use the launch in context parameters, definitions, and examples.

“Launch in context task parameters” on page 1153

Use the launch in context feature task parameters to perform tasks.

“Available windows” on page 1154

You can use windows with the inbound launch in context feature to start IBM Tivoli Storage Productivity Center. These windows are called with the `-function window name` argument.

Launch from the command line:

This topic describes how to start windows and tasks from the Windows AIX, Linux, or UNIX command line.

The bat file `tpc.bat` is used for Windows operating systems. The `.sh` file `TPCD.sh` is used for AIX, Linux, or UNIX systems. The parameters follow the file name. Parameters are separated by spaces. Parameters are optional and can appear in any order. See “Launch in context parameters” on page 1151, “Launch in context task parameters” on page 1153, and “Available windows” on page 1154 for the parameters.

Frequently used commands can be saved as unique command files and rerun as needed.

1. In the Command Prompt window, type the command for the action that you want to take, and press Enter.
The action occurs on the remote Data server. The graphical user interface (GUI) window for that action opens on the local computer.
2. Enter parameters if you previously did not enter any parameters. If you do not provide a user name, password, or server name, the IBM Tivoli Storage Productivity Center GUI opens the connection dialog box. If you do not provide these values, the GUI presents an error message box and requests the values. If you specify incorrect parameters, you receive an error message.

The following example launches a task from the Windows command line:

```
tpc.bat -user admin -passwd password1 -server
tivoli122.storage.usca.ibm.com:9549 -function
data.filesystem_capacity_by_computer -computer_id <computerID>
```

The following example uses the launch in context feature to view a window on the remote machine:

```
tpc.bat -user admin -passwd password1 -server  
tivoli122.storage.usca.ibm.com:9549 -function disk.storage_subsystem
```

Restrictions:

- When you open a plan in the SAN Planner using the command line or a Web browser for launch in context, you must enter the exact name of the plan. The value for **-plan_id** in the command line is case-sensitive and must match the name of the plan when it was created. For example, if you created a plan named **UserPlan1**, you might enter the following command on the command line:

```
./TPCD.sh -user administrator -passwd passw0rd -server  
tpcserver14.storage.xyz.com:9549 -function disk.san_Planner -plan_id  
UserPlan1
```

- There is a known limitation when you start the Filesystem Capacity > By Computer report using launch in context. To start this report, you must specify the ID for the computer upon which you want to report. For example:

```
./TPCD.sh -user root -passwd "passw0rd" -server  
tpcserver14.storage.xyz.com:9549 -function  
data.filesystem_capacity_by_computer -computer_id tpcserver14
```

The value you enter for "computer_id" might not match the ID for that computer that is stored in the Tivoli Storage Productivity Center database repository. You will receive a warning message if the ID for the computer you enter in the command line does not match the ID stored in the data repository. To ensure that this report is displayed properly when run from the launch in context command line or through Java Web Start, make sure to enter an ID for the computer that matches the ID that is displayed for it through the Topology Viewer or appropriate report in the product's user interface.

Launch from Start menu:

This topic describes how to start windows and actions from the Windows Start menu.

The bat file, `tpc.bat`, is used for Windows operating systems. The parameters follow the bat file name. Parameters are optional and can appear in any order. To use the launch in context feature from the Windows Start menu, complete the following steps:

1. Click **Start > Programs > IBM Tivoli Storage Productivity Center**, right-click **Productivity Center** and select **Properties**. The Productivity Center Properties window opens.
2. Enter the command-line arguments in the **Target** text field and click **Enter**.
If the Data server is remote, the action occurs on the Data server and the graphical user interface (GUI) window the action opens is on the local computer.
3. Enter the parameters if you did not previously enter the necessary parameters. If you do not provide the user name, password, or server name, the Tivoli Storage Productivity Center GUI opens the connection dialog box. If you do not provide these values, the GUI presents an error message box and requests the values.

The following example starts a task from the Windows Start menu:

```
"C:\Program Files\IBM\TPC\gui\tpc.bat" -user admin -passwd password1  
-server tivoli122.storage.usca.ibm.com:9549 -function  
data.filesystem_capacity_by_computer -computer_id computerID
```

The following example uses the Windows Start menu to view a window on the remote server:

```
"C:\Program Files\IBM\TPC\gui\tpc.bat" -user admin -passwd password1  
-server tivoli122.storage.usca.ibm.com:9549 -function  
disk.storage_subsystem
```

For more information see:

“Launch in context parameters” on page 1151

You can use the launch in context parameters, definitions, and examples.

“Launch in context task parameters” on page 1153

Use the launch in context feature task parameters to perform tasks.

“Available windows” on page 1154

You can use windows with the inbound launch in context feature to start IBM Tivoli Storage Productivity Center. These windows are called with the

-function *window name* argument.

Launch from Tivoli Enterprise Portal:

This topic shows you how to start an existing launch definition from the Tivoli Enterprise Portal graphical user interface (GUI).

You must configure Tivoli Enterprise Portal before you can use the launch in context feature. For information on prerequisites and details on how to configure the launch in context feature with Tivoli Enterprise Portal, see “Configuring Tivoli Enterprise Portal.” After Tivoli Enterprise Portal has been configured, you can start the application from the configuration window or use **Launch** from the Tivoli Enterprise Portal navigation tree.

If you use the Tivoli Enterprise Portal navigation tree, **Launch** will open the Tivoli Enterprise Portal graphical user interface with your predefined launch in context commands. To use one of these commands, complete the following steps:

1. In the Tivoli Enterprise Portal window, expand **Enterprise > Windows Systems *your_server_name* > Windows OS**. Right-click **Disk** and select **Launch**. Disk is one of several values you can select. Values include Disk, Network, Enterprise Services, Memory, Printer, Process, Processor, and System.
2. The Create or Edit Launch Definitions window opens showing a list of defined launch definitions in the **Existing Launch Definitions** area.
3. Select the launch definition you want to use from the list and click **Launch**. The window or task that is defined by that launch definition opens.

For information on editing the launch in context command with the Tivoli Enterprise Portal GUI, see:

Related tasks:

“Editing launch definitions” on page 67

You can edit an existing launch definition with the Tivoli Enterprise Portal graphical user interface (GUI).

Configuring Tivoli Enterprise Portal:

You can configure Tivoli Enterprise Portal so that you can use the launch in context feature.

To configure Tivoli Enterprise Portal complete the following steps:

1. The initial steps depend on whether you are going to open a window or perform a task.
 - a. If you are opening a window, in the Tivoli Enterprise Portal navigation window on the local computer, select where you want the launch definition to originate from. This opens the corresponding view on the target computer. Right-click one of the following locations.
 - The navigator item
 - A row in a table or situation event console view
 - A slice of a pie chart
 - A bar of a bar chartIf the launch definition originates from a table, chart, or situation event console view, open the workspace containing the view. If you think you might edit the argument later, create the launch definition from a table or chart view.
 - b. If you are going to perform a task, create a launch definition for one of the following tasks:
 - Create a volume or VDisk
 - Open the SAN Planner
 - Create a capacity report
 - Create a wasted space report
2. In the Tivoli Enterprise Portal window, expand **Enterprise > Windows Systems > your_server_name > Windows OS**. Right-click **Disk** and select **Launch**. Disk is one of several values you can select. Values include Disk, Network, Enterprise Services, Memory, Printer, Process, Processor, and System.
3. The Create or Edit Launch Definitions window opens. Click **Create New**.
4. In the **Name** text field, type a descriptive name for this launch definition. This name is used in the definition list.
5. In the **Target** text box click **Browse** to locate the tpc.bat file. The Select files for directories for Launch Definitions window opens. A typical location for Windows systems is C:\Program Files\IBM\TPC\gui\tpc.bat. Select the correct directory and click **OK**.

When you start a launch definition from the command line, if any part of the path contains spaces, enclose the entire entry in quotation marks.

6. In the **Arguments** text field, enter any parameters that will be added when the application is started. Arguments are entered as *-parameter value*. In the Create or Edit Launch Definitions window, click **Browse** next to **Arguments** to locate the parameters that you want to use. The Select Browser Options window opens. Click **Substitutable Items** and then click **OK**. Select parameters from the **Selected Context** and **Extended Context** lists.

The **help** argument lists the available **function** values. You must type the help parameter by itself into the **Arguments** text field or into the form tpc.bat -help functions on the command line for Windows or TPCD.sh -help functions for UNIX. The **-help functions** parameter provides a window that lists the nodes in the navigation tree that are available to start. If you type **-help** without the **functions** parameter, it results in an error message.

When all parameters have been selected, click **OK** to return to the Create or Edit Launch Definitions window.

7. If the parameters can be changed at launch time (for example, you want to change the system ID), check the **Edit arguments before execution** box.
8. Specify a value for the **Start in** location. Click **Browse** to locate the IBM Tivoli Storage Productivity Center installation and select the folder **gui** where `tpc.bat` is located.
9. Click **Evaluate** to see how the command resolves when the application starts. Correct any of the parameters that are incorrect and reevaluate the command.
10. When you are satisfied with the launch definition, click **Apply** to save your changes. This keeps the window open so that you can start the current launch definition or create a new one. Click **OK** to save the launch definition and close the window.
11. To start the current definition, select **Launch**. The Tivoli Storage Productivity Center graphical user interface (GUI) with the function parameter that you have provided is opened.

Related reference:

“Launch in context parameters” on page 1151

You can use the launch in context parameters, definitions, and examples.

“Launch in context task parameters” on page 1153

Use the launch in context feature task parameters to perform tasks.

“Available windows” on page 1154

You can use windows with the inbound launch in context feature to start IBM Tivoli Storage Productivity Center. These windows are called with the `-function window name` argument.

Editing launch definitions:

You can edit an existing launch definition with the Tivoli Enterprise Portal graphical user interface (GUI).

If the **Edit arguments before execution** box was checked when a launch definition was configured, you can edit the launch definition by changing the attribute values. To edit a launch definition, complete the following steps:

1. In the Tivoli Enterprise Portal window, expand **Enterprise > Windows Systems > your_server_name > Windows OS**. Right-click **Disk** and select **Launch**. Disk is one of several nodes you can select. Nodes include Disk, Network, Enterprise Services, Memory, Printer, Process, Processor, and System.
2. The Create or Edit Launch Definitions window opens showing a list of launch definitions in the **Existing Launch Definitions** area.
3. Select the launch definition that you want to edit. The values are displayed.
4. Make the necessary changes in the text fields.
5. Click **Evaluate** to see how the arguments will resolve when the application starts. Correct any incorrect arguments and re-evaluate the command.
6. Click **Apply** to save the changes and keep the window open. Click **OK** to save the launch definition and close the window.

If you do not want to overwrite the existing launch definition, save the new launch definition with a unique name.

Outbound launch and launch in context features

With the outbound launch and launch in context features, you can start the administrative applications for devices such as storage systems and switches from various points in the graphical user interface (GUI).

There are two options for starting an application from Tivoli Storage Productivity Center. The option that is available depends on the device that you want to start.

The launch feature starts an application at the starting point for that application.







The launch in context feature starts an application at a functionally related location in that application. Launch in context is available for the following applications:




- Tivoli Storage Productivity Center for Replication
- SAN Volume Controller 6.1 or later
- Storwize V7000
- Storwize V7000 Unified

Launch menu

When you right-click one of the following entities in Tivoli Storage Productivity Center GUI, you see a launch menu for that entity. The content of the launch menu varies as described in “Launch menu items” on page 69.

- **Alert Log Viewer.** Each row of the alert log is enabled for a launch in context menu. The context is determined by the type of the alert. Alert types include:
 - Replication
 - Storage Subsystem
 - Computer
 - Fabric
 - Switch
 - Port
 - Tape Library
- **Disk Manager Storage Subsystems table**
- **Fabric Manager Fabrics table**
- **Tape Manager Tape Libraries table**
- **Element Manager page** (of the Configuration Utility)
- **Topology Viewer objects** (see the following table)

Topology Object	Icon
Fabric	
Switch	
Switch port	
Tape library	
Storage subsystem	
Storage subsystem port	

Topology Object	Icon
File system volume	
Storage disk	
Storage pool	

Launch menu items

The following are the possible menu items that can be displayed in the launch menu:

- **Launch in context defined application.** This menu item is displayed for each application that is defined by a launch descriptor file. Tivoli Storage Productivity Center provides launch descriptor files for the following applications:
 - DS8000 Storage Manager
 - The management console for the storage systems that are listed under **Management GUI**.
 - Tivoli Storage Productivity Center for Replication

If the corresponding launch descriptor files are found for the applications, the following menu items are displayed:

DS8000 Storage Manager

Click this menu item to start DS8000 Storage Manager.

Management GUI

Click this menu item to start the management console for the following storage systems:

- SAN Volume Controller 6.1 or later
- Storwize V7000
- Storwize V7000 Unified

Replication Storage Systems Overview

Click this menu item to start Tivoli Storage Productivity Center for Replication. The window that is opened in the Tivoli Storage Productivity Center for Replication GUI depends on the context for the entity that you selected. For example, if you right-click a storage device on the **Element Manager** tab of the Configuration Utility and click **Replication Storage Systems Overview**, the Tivoli Storage Productivity Center for Replication Storage Systems window opens.

If the launch descriptor file cannot be found, the **Launch file items error** menu item is displayed. Click this item to view a description of the problem. The **Launch file items error** menu item is also displayed in situations such as the IBM Tivoli Integrated Portal server is not installed or is not accessible, the Tivoli Storage Productivity Center Data server is not accessible, or another system error occurred.

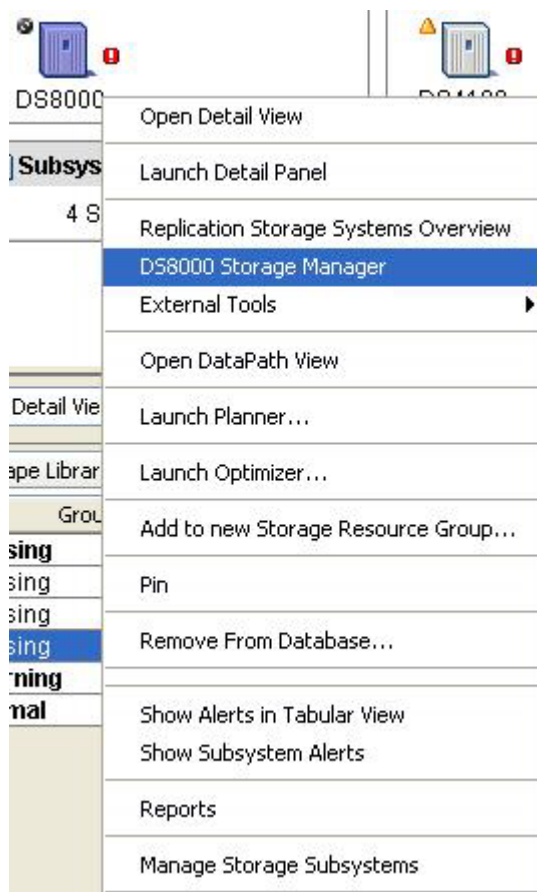
- **External Tools.** This menu item shows any external tools that are defined in the Configuration Utility. Examples of external tools include web-based user

interfaces and local executable files. For example, you might add the address of a website that you frequently access or an application that you often use such as a text editor or a spreadsheet.

- **Element Manager.** This menu item shows any element managers that were discovered for the device by a Storage Management Initiative Specification (SMI-S) query. To avoid menu item duplication, this item is not displayed when the associated element manager is defined by a launch descriptor file. In this situation, the element manager is displayed as a separate menu item.
- **User Defined Tools.** This menu item shows any user-defined element managers that are defined in the Configuration Utility. These element managers differ from the element managers in the **Element Manager** menu because they are user-defined and are not discovered by a SMI-S query.

Example launch menu

The following figure shows an example launch menu for a storage system in the **L0:Storage** topology view. In this example, the menu items **Replication Storage Systems Overview**, **DS8000 Storage Manager**, and **External Tools** are displayed. The **Element Manager** menu item is not displayed because no element managers other than **DS8000 Storage Manager** were discovered. The **User Defined Tools** menu item is not displayed because there are no user-defined element managers that are defined in the Configuration Utility.



Using the single sign-on feature with launch or launch in context

Single sign-on is an authentication process in which you can access more than one system or application by entering a single user ID and password. Single sign-on enables you to start applications from Tivoli Storage Productivity Center without providing logon credentials for the application. Single sign-on is available only for Tivoli Storage Productivity Center for Replication, DS8000, SAN Volume Controller 6.1 and later, and Storwize V7000.

For Tivoli Storage Productivity Center for Replication Storage Manager, single sign-on is enabled by default. For DS8000 Storage Manager, SAN Volume Controller, and Storwize V7000, single sign-on must be enabled.

To enable single sign-on, see the following topics in the Tivoli Storage Productivity Center information center:

- Configuring Tivoli Storage Productivity Center for DS8000 LDAP authentication
- Configuring Tivoli Storage Productivity Center for SAN Volume Controller and Storwize V7000 single sign-on

Related concepts:

“Single sign-on” on page 58

Single sign-on is an authentication process that enables you to enter one user ID and password to access multiple applications. For example, you can access IBM Tivoli Integrated Portal and then access IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication from Tivoli Integrated Portal using a single user ID and password. Single sign-on integrates with the launch and launch in context features to enable you to move smoothly from one application to another application.

User roles

IBM Tivoli Storage Productivity Center provides user roles that can be mapped to user groups contained in the authentication repository, which is either the local operating system or an LDAP-compliant directory, depending on the choice of authentication mechanism made during the Tivoli Storage Productivity Center installation. The roles determine the user's ability to see and act upon specific nodes in the navigation tree of the GUI. Other factors that affect the user's capability are the components that are licensed, installed, and running.

The superuser is the highest level role and can be used to install IBM Tivoli Storage Productivity Center and to initially set up the product. The superuser has authority over all the other roles in Tivoli Storage Productivity Center:

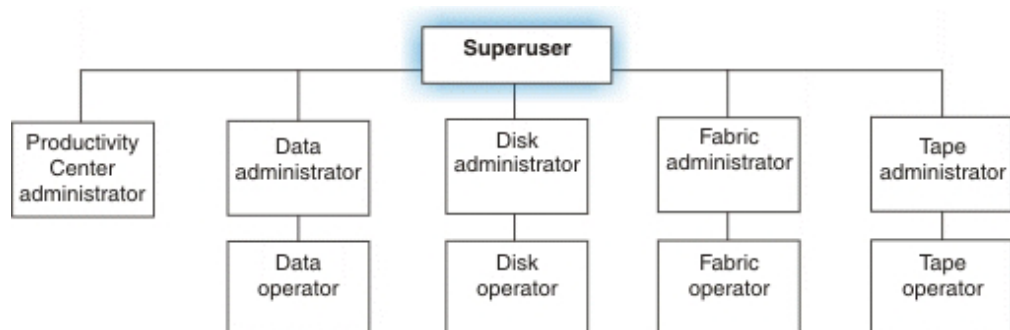


Figure 3. User roles in Tivoli Storage Productivity Center

The functions that are visible for the user roles require that the following components be licensed and installed as shown in Table 6.

Table 6. Functions that require components to be licensed and installed

Function	Component
Data Manager	IBM Tivoli Storage Productivity Center for Data
Disk Manager (requires Device server installation)	IBM Tivoli Storage Productivity Center for Data or Tivoli Storage Productivity Center for Disk or IBM Tivoli Storage Productivity Center for Fabric
Disk performance function (requires Device server installation)	Tivoli Storage Productivity Center for Disk
Fabric Manager (requires Device server installation)	IBM Tivoli Storage Productivity Center for Fabric or IBM Tivoli Storage Productivity Center for Data or Tivoli Storage Productivity Center for Disk
Fabric performance function (requires Device server installation)	IBM Tivoli Storage Productivity Center for Fabric
Tape Manager (requires Device server installation)	IBM Tivoli Storage Productivity Center for Fabric or IBM Tivoli Storage Productivity Center for Data or Tivoli Storage Productivity Center for Disk
Element Manager	IBM Tivoli Storage Productivity Center Basic Edition or IBM Tivoli Storage Productivity Center Standard Edition

Table 7 shows the GUI nodes that require special roles to view, edit, and act upon the services. The Productivity Center administrator or superuser is required for viewing configurations.

Table 7. GUI nodes that require special roles or licenses

GUI node	Special roles
Administrative Services	All roles can see this node but the user must be a Productivity Center administrator or superuser to edit and act upon the services (for example enable or disable tracing, shut down services, delete services, and so forth).
Administrative Services - Configuration - Role-to-Group Mappings	Only the Productivity Center administrator or superuser can see this node.
Administrative Services - Configuration - License Keys	Only the Productivity Center administrator or superuser can see this node.
Tivoli Storage Productivity Center - My Reports - System Reports - Data	Requires the data administrator role to create or edit objects, or to run jobs, (for example, pings and scans). Requires the data operator role to view reports.
Tivoli Storage Productivity Center - My Reports - System Reports - Disk	Requires that the disk performance function be installed and the disk operator role.

Table 7. GUI nodes that require special roles or licenses (continued)

GUI node	Special roles
Tivoli Storage Productivity Center - My Reports - System Reports - Fabric	Requires the fabric administrator role to create or edit objects, or to run jobs (for example, pings and scans). Requires the fabric operator role to view reports. This includes being able to view the fabric assets, port connections, SAN switch report, switch port errors, and switch ports (packet rate and data rate).
Tivoli Storage Productivity Center - My Reports - Batch Reports	Requires any administrator role to create the batch reports.
Tivoli Storage Productivity Center - Monitoring - Probes	Requires any administrator role to create or edit probes.
Tivoli Storage Productivity Center - Alerting - Alert log	Requires any administrator role to delete or clear alerts. The Storage subsystems node appears if any disk array is installed.
Data Manager	Requires the data administrator role to create or edit objects, or to run jobs (for example, pings and scans). Requires the data operator role to view reports.
Data Manager for Databases	Requires the data administrator role to create or edit objects, or to run jobs (for example, pings and scans). Requires the data operator role to view reports.
Data Manager for Chargeback	Requires the data administrator or data operator role.
Disk Manager	Requires the disk administrator role to create or edit objects, or to run jobs (for example, pings and scans). Requires the disk operator role to view reports. The Storage Subsystem Performance Monitors node requires that the Disk performance function is available.
Disk Manager - Monitoring - Subsystem Performance Monitors - Profile Management - Reporting - Storage Subsystem Performance	Disk performance function must be available.
Fabric Manager	Requires the fabric administrator role to create or edit objects, or to run jobs (for example, pings and scans). Requires the fabric operator role to view reports.
Fabric Manager - Monitoring - Switch Performance Monitors - Reporting - Switch Performance - Constraint Violations	Fabric performance function must be available.

Table 7. GUI nodes that require special roles or licenses (continued)

GUI node	Special roles
Tape Manager	Requires the tape administrator role to create or edit objects, or to run jobs (for example, pings and scans). Requires the tape operator role to view reports.
Element Manager	All roles can view and edit objects and run jobs.

User tasks

This section provides an overview of the key IBM Tivoli Storage Productivity Center tasks that you can perform. It includes tasks that are common to all configurations and tasks that require that specific components be installed and licensed.

Common tasks

Perform common tasks on any Tivoli Storage Productivity Center configuration. These tasks are available under the **Administrative Services** and **IBM Tivoli Storage Productivity Center** nodes in the navigation tree.

- View and manage the services for the Data server, Device server, and Replication server.
- Configure devices for monitoring by Tivoli Storage Productivity Center.
- Manage agents and data sources, including CIM agents, Storage Resource agents, Data agents, Fabric agents, storage systems, Tivoli Storage Productivity Center servers, and VMware VI.
- Discover storage devices that are visible to Tivoli Storage Productivity Center.
- Configure and customize Tivoli Storage Productivity Center for your environment using the functions available under the **Administrative Services > Configuration** node in the navigation tree.
- Authorize users by associating user groups with roles (requires superuser role).
- Deploy Storage Resource agents on hosts that you want to monitor.
- Manage all the schedules, runs, and jobs that are related to the monitoring of storage devices by Tivoli Storage Productivity Center.
- Define probe schedules to collect information about storage devices and other Tivoli Storage Productivity Center servers.
- Create batch reports for information that you want to view on a regular basis in an output file.
- View System Reports that contain key information about the devices that are monitored by Tivoli Storage Productivity Center. These reports are automatically generated when data collection schedules are run.
- Generate Rollup Reports to view information that has been collected by Tivoli Storage Productivity Center servers if you have multiple servers installed in your environment.
- Create and manage groups of related storage entities. Use storage resource groups as input into the SAN Planner.
- Analyze storage subsystems, identify performance bottlenecks, and recommend changes to improve performance (Storage Optimizer).
- Create end-to-end planning for fabrics, hosts, storage controllers, storage pools, volumes, paths, ports, zones, zone sets, and replication (SAN Planner).

- Determine whether existing SAN configurations violate defined best practices or policies (Configuration Analysis).
- Check the alert log to see detailed information about alerts.
- Access a topology view of your storage infrastructure.
- Use Tape Manager to view tape libraries, create tape library groups and run probes against those groups, and view tape library alerts.

Data Manager tasks

Data Manager tasks are tasks that are specific to the Data Manager, Data Manager for Databases, and Data Manager for Chargeback components.

The following tasks are available under the **Data Manager**, **Data Manager for Databases**, and **Data Manager for Chargeback** nodes in the navigation tree:

- Create monitoring groups that contain related storage devices. Use groups when you want to perform the same action against multiple devices, such as collecting data with a scan or probe. In Data Manager, you can create groups of computers, file systems, directories, users, OS user groups, databases, tablespaces, and tables.
- Define ping schedules to collect availability information about computers and clusters.
- Define scan schedules to collect detailed information about storage usage and trending for computers, file systems, directories, and clusters.
- Create profiles to define the files that are scanned during a file systems scan and the statistics that are gathered.
- Create and manage alerts for computers, file systems, directories, database instances, tablespaces, and tables. Use alerts to be notified of conditions that are detected on storage devices.
- Establish quotas and constraints to define storage use limits for users and user groups.
- Automatically extend local file systems of managed hosts.
- Generate and view data system reports.
- Use Data Manager for Chargeback to generate invoices that charge back for storage usage.

Disk Manager tasks

The Disk Manager tasks are tasks that are specific to the Disk Manager component.

The following tasks are available under the **Disk Manager** node in the navigation tree:

- Configure storage systems for monitoring by Tivoli Storage Productivity Center.
- Create and manage virtual disks and volumes.
- Define and launch element managers for storage systems.
- Use the Storage Optimizer to identify performance bottlenecks and recommend changes to improve performance for storage systems.
- Use the SAN Planner to perform end-to-end planning for fabrics, hosts, storage controllers, storage pools, volumes, paths, ports, zones, zone sets, and replication.
- Create monitoring groups that contain related storage systems. Use these groups when you want to perform the same action against multiple storage systems, such as collecting data with a probe.

- Define performance monitors to collect performance data related to storage systems.
- Create and manage alerts for storage systems. Use alerts to be notified of conditions that are detected on storage systems.
- Create workload profiles to be used by the SAN Planner to describe characteristics of the volumes to be created.
- Create provisioning profiles to be used by the SAN Planner to define the actual capacity (and security - zoning) characteristics of whatever the profile happens to be associated with.
- Create thin provisioning profiles to specify the parameters for thin provisioned volumes to be created through SAN Planner.
- Generate and view detailed reports about the storage systems that are monitored by Tivoli Storage Productivity Center.

Fabric Manager tasks

Fabric Manager tasks are tasks that are specific to the Fabric Manager component.

The following tasks are available under the **Fabric Manager** node in the navigation tree:

- Configure fabrics and switches for monitoring by Tivoli Storage Productivity Center.
- Create and work with fabric zones, zone aliases, and zone sets
- Create monitoring groups that contain related fabrics. Use these groups when you want to perform the same action against multiple fabrics, such as collecting data with a probe.
- Define switch performance monitors to collect performance data related to switches.
- Create and manage alerts for fabrics, switches, and endpoint devices. Use alerts to be notified of conditions that are detected on these storage devices.
- Generate and view fabric system reports and switch performance reports.

Tivoli Storage Productivity Center for Replication tasks

You can use IBM Tivoli Storage Productivity Center for Replication to obtain continuous availability and disaster recovery solutions by using point-in-time replication, which includes FlashCopy, and Snapshot, and continuous replication, which includes Metro Mirror and Global Mirror.

You can set up Metro Mirror and Global Mirror sessions to replicate data in both the forward and reverse directions.

Use Tivoli Storage Productivity Center for Replication for the following tasks:

- Manage Tivoli Storage Productivity Center for Replication servers
- Create and monitor storage system connections
- Create and monitor logical paths
- Create and manage sessions for performing these types of data replication.
 - FlashCopy (ESS, DS6000, DS8000, and SAN Volume Controller)
 - Snapshot (the XIV system)
 - Global Mirror Single Direction (ESS, DS6000, DS8000, and SAN Volume Controller)
 - Global Mirror Either Direction with Two-Site Practice (ESS, DS6000, and DS8000)

- Global Mirror Failover/Failback (ESS, DS6000, DS8000, SAN Volume Controller, and the XIV system)
- Global Mirror Failover/Failback with Practice (ESS, DS6000, DS8000, and SAN Volume Controller)
- Metro Mirror Single Direction (SAN Volume Controller)
- Metro Mirror Failover/Failback (SAN Volume Controller, and the XIV system)
- Metro Mirror Failover/Failback with Practice (SAN Volume Controller)
- Metro Mirror or Global Copy (ESS, DS6000, and DS8000)
- Metro Mirror Failover/Failback or Global Copy (ESS, DS6000, and DS8000)
- Metro Mirror Failover/Failback with HyperSwap (ESS, DS6000, and DS8000)
- Metro Global Mirror Failover/Failback (DS8000)
- Metro Global Mirror Failover/Failback (ESS and DS8000)
- Metro Global Mirror Failover/Failback with Practice (ESS and DS8000)
- Metro Global Mirror with HyperSwap (ESS and DS8000)

Note: Metro Global Mirror is available only with IBM Tivoli Storage Productivity Center for Replication Three Site Business Continuity. Tivoli Storage Productivity Center for Replication Three Site Business Continuity requires a license key.

- Practicing recovery procedures while maintaining disaster recover capability

Data collection

Determining what information you want to gather about the storage resources within your enterprise is critical to helping you implement a storage management strategy. Create data collection schedules to gather information about the storage resources in your environment. There are different data collection schedules within Tivoli Storage Productivity Center depending on the type of information that you want to gather and from which types of resources the information is retrieved: discovery, probes, scans, pings, and performance monitors.

You must run discovery schedules or use the Configure Devices wizard before you can collect more detailed information about storage resources. After a storage resource is discovered by Tivoli Storage Productivity Center, use probes, scans, pings, and performance monitors to collect detailed information about those resources. You do not need to run discovery schedules for devices that use native API connection.

The data and statistics gathered by data collection schedules are stored in the database repository. You must define and run schedules before you can perform the other tasks within Tivoli Storage Productivity Center, such as viewing reports, enforcing storage policies, setting alerts, using the topology viewer, and managing storage subsystems.

The following types of data collection schedules are available in Tivoli Storage Productivity Center:

Availability Monitoring

Use pings to collect information about the availability of the storage assets in your enterprise. You can run pings against servers and computers only. Pings are unavailable for storage subsystems and other devices.

Asset Collection

Use probes to collect statistics about the storage assets in your enterprise,

such as computers, disk controllers, hard disks, file systems, storage subsystems, fabrics, switches, hypervisors, and RDBMS objects.

Usage Monitoring

Use scans to collect statistics about the usage and trends of your actual storage consumption. Scans provide all the data for usage reporting, quota analysis, and constraint analysis. They gather information at the file and directory level. You can run scans against servers and computers only. Scans are unavailable for storage subsystems, switches, fabrics, and other devices.

Switch Performance Monitoring

Use fabric probes and switch performance monitors to collect information about switches and their performance.

Storage Subsystems Performance Monitoring

Use storage subsystem probes and subsystem performance monitors to collect information about storage subsystems and their performance.

Tivoli Storage Productivity Center Server Monitoring

Use TPC Server probes to collect the storage asset and capacity information that has been gathered by other Tivoli Storage Productivity Center servers in your environment. Create master and subordinate relationships among servers to use a single interface to generate reports based on data and metrics collected by multiple servers.

You can create groups and profiles to use with data collection schedules to more efficiently monitor storage resources.

Monitoring Groups

Use monitoring groups to ping, probe, or scan multiple storage resources using the same schedule. You can create named groups of computers, storage subsystems, file systems, directories, fabrics, databases, table spaces, tables, users, and OS user groups.

For example, to gather information about all of the computers in a payroll department, create a computer monitoring group named Payroll and populate that group with the computers from the payroll department. Then, to collect information about those computers, select the computer monitoring group named Payroll when defining a data collection schedule. As computers are added or removed from the payroll department, update the Payroll computer monitoring group with those changes. Any data collection schedule that uses Payroll automatically collects information about the new computers in the group during its next run.

Profiles

Use profiles to fine-tune and control the files to be scanned during a file system scan and to specify the statistical information to be gathered.

Default profiles are supplied with Data Manager with Data Manager for Databases. With the default profiles, you can select a specific statistic that you want to gather. If you want to scan a computer or file system and gather all the available statistical information, select all of the default profiles when you create your scan.

You can also use profiles with **Reporting > Usage** reports to determine what information is displayed in those reports.

Data replication

Tivoli Storage Productivity Center for Replication supports several methods of data replication. The type of data replication that is associated with a given session is known as the *session type* (also known as a copy type).

The following table describes the session types that are supported by Tivoli Storage Productivity Center for Replication. Depending on the edition of Tivoli Storage Productivity Center for Replication that you are using, some of these session types might not be available.

Table 8. Session type summary

Copy type	Supported Software	Supported storage systems	Description
Basic HyperSwap	Tivoli Storage Productivity Center for Replication Basic Edition for System z and Tivoli Storage Productivity Center for Replication for System z	<ul style="list-style-type: none"> TotalStorage Enterprise Storage Server Model 800 System Storage DS8000 System Storage DS6000 	Basic HyperSwap replication is a special Metro Mirror replication method designed to provide high availability in the case of a disk storage system failure. Using Basic HyperSwap with Metro Mirror, you can configure and manage your synchronous Peer-to-Peer Remote Copy (PPRC) pairs.
FlashCopy	Tivoli Storage Productivity Center all editions	<ul style="list-style-type: none"> TotalStorage Enterprise Storage Server Model 800 System Storage DS8000 System Storage DS6000 SAN Volume Controller Storwize V7000 Storwize V7000 Unified 	<p>FlashCopy replication creates a point-in-time copy in which the target volume contains a copy of the data that was on the source volume when the FlashCopy was established. Using FlashCopy, your data exists on the second set of volumes in the same storage system, and can be restored to the first set of volumes.</p> <p>SAN Volume Controller or Storwize V7000 FlashCopy sessions are managed by using FlashCopy consistency groups. Sessions for IBM TotalStorage Enterprise Storage Server (ESS) and IBM DS6000 and DS8000 are not managed by using FlashCopy consistency groups.</p>

Table 8. Session type summary (continued)

Copy type	Supported Software	Supported storage systems	Description
Snapshot	Tivoli Storage Productivity Center all editions	The XIV system	Snapshot is a session type that creates a point-in-time copy of a volume or set of volumes without having to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.
Metro Mirror Single Direction	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified 	Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. The source is located in one storage system and the target is located in another storage system. Using Metro Mirror, your data exists on the second site that is less than 300 KM away, and can be restored to the first site.
Metro Mirror Failover/Failback	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified • The XIV system 	<p>Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. Using Metro Mirror Failover / Failback, your data exists on the second site that is less than 300 KM away. You can use failover and failback to switch the direction of the data flow. This ability enables you to run your business from the secondary site.</p> <p>Using Metro Mirror with HyperSwap, your data exists on the second site that is less than 300 KM away. The data can be restored to the first site. You can also use failover for a backup copy of the data if your primary volumes encounter a permanent I/O error.</p>

Table 8. Session type summary (continued)

Copy type	Supported Software	Supported storage systems	Description
Metro Mirror Failover/Failback with Practice	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified 	Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. The source is located in one storage system and the target is located in another storage system. Metro Mirror Failover / Failback with Practice combines Metro Mirror and FlashCopy to provide a point-in-time copy of the data on the remote site.
Global Mirror Single Direction	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified 	Global Mirror is a method of asynchronous, remote data replication between two sites that are over 300 KM apart. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system. Using Global Mirror, your data exists on the second site that is more than 300 KM away, and can be restored to the first site.
Global Mirror Either Direction with Two-Site Practice	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 	Global Mirror is a method of asynchronous, remote data replication between two sites that are over 300 KM apart. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system. Global Mirror Either Direction with Two-Site Practice combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on either the primary or secondary sites that are over 300 KM apart.

Table 8. Session type summary (continued)

Copy type	Supported Software	Supported storage systems	Description
Global Mirror Failover/Failback	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified • The XIV system 	<p>Global Mirror is a method of asynchronous, remote data replication between two sites that are over 300 KM apart. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system. Using Global Mirror Failover / Failback, your data exists on the second site that is more than 300 KM away, and you can use failover and failback to switch the direction of the data flow. This ability enables you to run your business from the secondary site.</p>
Global Mirror Failover/Failback with Practice	Tivoli Storage Productivity Center for Replication Two Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 • System Storage DS8000 • System Storage DS6000 • SAN Volume Controller • Storwize V7000 • Storwize V7000 Unified 	<p>Global Mirror is a method of asynchronous, remote data replication between two sites that are over 300 KM apart. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system. Global Mirror Failover / Failback with Practice combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on a remote site at a distance over 300 KM away from your first site.</p>

Table 8. Session type summary (continued)

Copy type	Supported Software	Supported storage systems	Description
Metro Global Mirror	Tivoli Storage Productivity Center for Replication Three Site Business Continuity	<ul style="list-style-type: none"> • TotalStorage Enterprise Storage Server Model 800 (only H1 site) • System Storage DS8000 	<p>Metro Global Mirror is a method of continuous, remote data replication that operates between three sites that varying distances apart. Metro Global Mirror combines Metro Mirror synchronous copy and Global Mirror asynchronous copy into a single session, where the Metro Mirror target is the Global Mirror source. Using Metro Global Mirror and Metro Global Mirror with HyperSwap, your data exists on a second site that is less than 300 KM away, and a third site that is more than 300 KM away. Metro Global Mirror uses both Metro Mirror and Global Mirror Failover / Failback to switch the direction of the data flow. This ability enables you to run your business from the secondary or tertiary sites.</p> <p>Using Basic HyperSwap with Metro Global Mirror, you can configure and manage three-site continuous replication needed in a disaster recovery event.</p>

Table 8. Session type summary (continued)

Copy type	Supported Software	Supported storage systems	Description
Metro Global Mirror with Practice	Tivoli Storage Productivity Center for Replication Three Site Business Continuity	<ul style="list-style-type: none"> TotalStorage Enterprise Storage Server Model 800 (only H1 site) System Storage DS8000 	Using Metro Global Mirror with Practice, you can practice your disaster recovery actions while maintaining disaster recovery capabilities. Your data exists on a second site that is less than 300 KM away, and a third site that is more than 300 KM away. Metro Global Mirror uses both Metro Mirror and Global Mirror Failover / Failback to switch the direction of the data flow; as a result, you can run your business from the secondary or tertiary sites, and simulate a disaster.

Basic HyperSwap (ESS, DS6000, and DS8000)

Basic HyperSwap is an entitled copy services solution for z/OS® version 1.9 and later. It provides high availability of data in the case of a disk storage system failure. Basic HyperSwap is not a disaster recovery solution. If a session is suspended but the suspend was not caused by a HyperSwap trigger, no freeze is done to ensure consistency of the session.

When HyperSwap is combined with Metro Mirror and Metro Global Mirror replication, you can prepare for disaster recovery and ensure high availability. If a session is suspended but the suspend was not caused by a HyperSwap trigger, a freeze is done to ensure consistency of the session.

Note: This replication method is available on only ESS, DS6000, and DS8000 storage systems, and on management servers running IBM Tivoli Storage Productivity Center for Replication Basic Edition for System z or IBM Tivoli Storage Productivity Center for Replication for System z.

Basic HyperSwap replication performs the following actions:

- Manages CKD volumes in Metro Mirror (synchronous peer-to-peer remote copy [PPRC]) relationships.
- Permits only CKD volumes to be added to the HyperSwap session. The graphical user interface (GUI) shows only CKD volumes when you add a copy set. The command-line interface (CLI) fails to add a copy set if a fixed block volume is specified.
- Monitoring for events that indicate a storage device has failed.
- Determining whether the failing storage device is part of a Metro Mirror (synchronous PPRC) pair.
- Determining the action to be taken from policy.
- Ensuring that data consistency is not violated.

- Swapping the I/O between the primary logical devices in the consistency group with the secondary logical devices in the consistency group. A swap can be from the preferred logical devices to the alternate logical devices or from the alternate logical devices to the preferred logical devices.

Metro Mirror Failover/Failback with HyperSwap

Metro Mirror Failover/Failback uses HyperSwap to configure and manage synchronous Peer-to-Peer Remote Copy (PPRC) pairs.

Metro Global Mirror with HyperSwap

Metro Global Mirror with HyperSwap is a z/OS replication feature that provides the three-site continuous replication needed in a disaster recovery event.

Important: If HyperSwap occurs by event when running a Metro Global Mirror with a HyperSwap session, a full copy of the data occurs to return to a full three-site configuration. If you issue a HyperSwap command when running a Metro Global Mirror with a HyperSwap session, a full copy does not occur. A full copy is required only for an unplanned HyperSwap or a HyperSwap initiated using the z/OS SETHS SWAP command.

Example

Jane is using multiple DS8000 storage systems. Her host applications run on z/OS and her z/OS environment has connectivity to the DS8000 storage systems. She has a site in Manhattan and a bunker in Hoboken. While she does not need a disaster recovery solution, she does need a high-availability solution to keep her applications running around the clock. Jane is worried that if a volume fails on her DS8000 in Manhattan, her database application will not be able to operate. Even a small downtime can cost Jane thousands of dollars. Jane uses a Basic HyperSwap session to mirror her data in Manhattan to her secondary DS8000 in Hoboken. If a volume at the Manhattan site fails, Basic HyperSwap automatically directs application I/O to the mirrored volumes in Hoboken.

FlashCopy


FlashCopy replication creates a point-in-time copy in which the target volume contains a copy of the data that was on the source volume when the FlashCopy was established.

The ESS, DS6000, and DS8000 platforms provide multiple logical subsystems (LSSs) within a single physical subsystem, while the following platforms provide multiple I/O groups:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

All platforms support local replication in which the source volume is located in one LSS or I/O group and the target volume is located in the same or another LSS or I/O group. Using FlashCopy, you can reference and update the source volume and target volume independently.

The following figure illustrates how a FlashCopy session works.



Example

Jane uses FlashCopy to make a point-in-time copy of the customer data in existing international accounts. Every night, the bank's servers perform batch processing. Jane uses FlashCopy to create checkpoint restarts for the batch processing in case the batch processing fails. In her batch processing, the first step is to balance all international accounts, with a FlashCopy taken of the resulting data. The second step in her batch processing is to process the international disbursements. If the second step in the batch process fails, Jane can use the FlashCopy made of the first step to repeat the second step, instead of beginning the entire process again. Jane also writes a CLI script that performs a FlashCopy every night at 11:59 PM, and another script that quiesces the database. She backs this data up on tape on her target storage system, and then sends the tape to the bank's data facility in Oregon for storage.

Snapshot

Snapshot is a session type that creates a point-in-time copy of a volume or set of volumes. You do not have to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.

The XIV system uses advanced snapshot architecture to create a large number of volume copies without affecting performance. By using the snapshot function to create a point-in-time copy, and to manage the copy, you can save storage. With the XIV system snapshots, no storage capacity is used by the snapshot until the source volume (or the snapshot) is changed.

The following figure illustrates how a Snapshot session works.



Example

Jane's host applications are using an XIV system for their back-end storage. With the XIV system, Jane can create a large number of point-in-time copies of her data. The snapshot function ensures that if data becomes corrupted, she can restore the data to any number of points in time.

Jane sets up a Snapshot session by using Tivoli Storage Productivity Center for Replication and specifies the volumes on the XIV system that are used by her host applications. Jane does not have to provision target volumes for all the snapshots she intends to make. She can quickly get a single Snapshot session configured and ready.

When the session is configured, Jane writes a CLI script that performs a Create Snapshot command to the session every two hours. If a problem occurs, such as data becoming corrupted, Jane can find a snapshot of the data before the problem occurred. She can restore the data to that point.

By creating a set of snapshots of the data, Jane can also schedule batch processing against that data every day. She can use the batch processing to analyze certain trends in the market without causing any effect to the host applications.

Metro Mirror

Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. The source is located in one storage system and the target is located in another storage system.

Attention: For Tivoli Storage Productivity Center for Replication for System z sessions containing Metro Mirror relationships, ensure that the session does not contain system volumes (such as paging volumes) unless the session is enabled for HyperSwap. If HyperSwap is not enabled, a freeze that is issued through Tivoli Storage Productivity Center for Replication might cause Tivoli Storage Productivity Center for Replication processing to freeze. This situation might prevent the session from ensuring that the data is consistent.


Metro Mirror replication maintains identical data in both the source and target. When a write is issued to the source copy, the changes made to the source data are propagated to the target before the write finishes posting. If the storage system goes down, Metro Mirror provides zero loss if data must be used from the recovery site.

A Metro Mirror session in Global Copy mode creates an asynchronous relationship to accommodate the high volume of data migration. As a result, the data on the target system might no longer be consistent with the source system. The Metro Mirror session switches back to a synchronous relationship when a Metro Mirror Start command is reissued. In addition, you can start a Metro Mirror session in Global Copy mode and toggle between Metro Mirror and Global Copy modes to accommodate time periods in which you value host input/output (I/O) response time over data consistency.

Tip: To determine if there is any out of synch data that must to be copied over before the session is consistent, check the percent that is complete in the session details panel.

Metro Mirror Single Direction


The following figure illustrates how a Metro Mirror Single Direction session works.



Metro Mirror Failover/Failback

Using Metro Mirror with failover/failback, your data exists on the second site that is less than 300 KM away, and you can use failover/failback to switch the direction of the data flow. This session type enables you to run your business from the secondary site, and to copy changes made at the second site back to the primary site when you want to resume production at the primary site.

The following figure illustrates how a Metro Mirror with Failover/Failback session works.




Metro Mirror Failover/Failback with Practice

A Metro Mirror Failover/Failback with Practice session combines Metro Mirror and FlashCopy to provide a point-in-time copy of the data on the remote site. You can use this session type to practice what you might do if a disaster occurred, without losing your disaster recovery capability.

This solution consists of two host volumes and an intermediate volume.

The following figure illustrates how a Metro Mirror Failover/Failback with Practice session works.



Metro Mirror Failover/Failback with HyperSwap

A Metro Mirror Failover/Failback session can be enabled to have HyperSwap capabilities. To enable HyperSwap the following circumstances must apply:

- The session is running on a Tivoli Storage Productivity Center for Replication server running on IBM z/OS.
- The volumes are only for TotalStorage Enterprise Storage Server, System Storage DS8000, and DS6000 systems.
- The volumes are count key data (CKD) volumes that are attached to the z/OS system.


Metro Mirror Failover/Failback with HyperSwap combines the high availability of Basic HyperSwap with the redundancy of a two-site Metro Mirror

Failover/Failback solution when managing count key data (CKD) volumes on z/OS. If the primary volumes encounter a permanent I/O error, the I/O is automatically swapped to the secondary site with little to no impact on the application.

A swap can be planned or unplanned. A planned swap occurs when you issue a HyperSwap command from the **Select Action** list in the graphical user interface (GUI) or when you issue a **cmdsess -action hyperswap** command.

The following figure illustrates how a Metro Mirror Failover/Failback session enabled for HyperSwap works.

Metro Mirror Failover/Failback



For more information about enabling HyperSwap, see “Managing a session with HyperSwap and Open HyperSwap replication” on page 100.

Metro Mirror Failover/Failback with Open HyperSwap

A Metro Mirror Failover/Failback session can be enabled to have Open HyperSwap capabilities. To enable Open HyperSwap the following circumstances must apply:


- The volumes in the session are System Storage DS8000 5.1 or later volumes.
- The volumes in the session are fixed block and are mounted to IBM AIX® 5.3 or AIX 6.1 hosts with the following modules installed:
 - Subsystem Device Driver Path Control Module (SDDPCM) version 3.0.0.0 or later
 - Multi-Path Input/Output (MPIO) module (the version that is provided with AIX version 5.3 or 6.1)
- The connections between the AIX host systems and the Tivoli Storage Productivity Center for Replication server have been established.

Metro Mirror Failover/Failback with Open HyperSwap combines the high availability of Basic HyperSwap on z/OS for fixed block AIX volumes with the redundancy of a two-site Metro Mirror Failover/Failback solution. If the primary volumes encounter a permanent I/O error, the I/O is automatically swapped to the secondary site with little to no impact on the application.

A swap can be planned or unplanned. A planned swap occurs when you issue a HyperSwap command from the **Select Action** list in the GUI or when you issue a **cmdsess -action hyperswap** command.

The following figure illustrates how a Metro Mirror Failover/Failback session enabled for Open HyperSwap works.

Metro Mirror Failover/Failback



For more information about enabling Open HyperSwap, see “Managing a session with HyperSwap and Open HyperSwap replication” on page 100.

Examples

Metro Mirror Single Direction

At the beginning of a work week, Jane is notified that between 10:00 AM and 11:00 AM on the next Friday, power in her building is going to be shut off. Jane does not want to lose any transactions during the power outage, so she decides to transfer operations to the backup site during the outage. She wants a synchronous copy method with no data loss for her critical business functions, so she chooses Metro Mirror, which can be used between locations that are less than 300 KM apart.

In a synchronous copy method, when a write is issued to change the source, the change is propagated to the target before the write is completely posted. This method of replication maintains identical data in both the source and target. The advantage of this is when a disaster occurs, there is no data loss at the recovery site because both writes must complete before signaling completion of a write to the source application. Because the data must be copied to both System Storage DS8000 devices before the write is completed, Jane can be sure that her data is safe.

The night before the planned outage, Jane quiesces her database and servers in San Francisco and starts the database and servers in Oakland. To accomplish this task, Jane issues the Suspend and Recover commands, and then issues the Start command on the secondary site. She powers off her equipment in San Francisco to avoid any power spikes during reboot after the power is turned back on.

Metro Mirror in Global Copy mode

At the beginning of a work week, Jane is notified that between 10:00 AM and 11:00 AM on the next Friday, power in her building is going to be shut off. Jane does not want to lose any transactions during the power outage, so she decides to transfer operations to their backup site during the outage. She wants a synchronous copy method with no data loss for her critical business functions, so she chooses Metro Mirror, which can be used between locations that are less than 300 KM apart.

Jane wants to limit her application impact while completing the initial Metro Mirror synchronization, so she begins her session in Global Copy mode. After she sees that about 70% of the data has been copied, Jane decides to switch the session into Metro Mirror mode, assuring data consistency.

Metro Mirror with Practice

Jane wants to run a Metro Mirror with Practice from San Francisco to Oakland. She wants to verify her recovery procedure for the Oakland site,

but she cannot afford to stop running her Metro Mirror session while she takes time to practice a recovery. By using a Metro Mirror with Practice session, Jane is able to practice her disaster recovery scenario in Oakland while her Metro Mirror session runs uninterrupted. By practicing running her applications at the Oakland site, Jane is better prepared to make a recovery if a disaster ever strikes the San Francisco site.

While her session is running in a Prepared state, Jane practices a recovery at her Oakland site by issuing the Flash command. This command momentarily pauses the session and starts a FlashCopy to the H2 volumes. As soon as the FlashCopy is started, her session will be restarted. These FlashCopy files create a consistent version of the data on the H2 volume that she can use for recovery testing, while her session continues to replicate data from San Francisco to Oakland. As a result, she can carry out her recovery testing without stopping her replication for any extended duration of time.

If at some point in time, the Metro Mirror session suspends due to a failure, Jane can use the practice session to restart her data replication while maintaining a consistent copy of the data at the Oakland site, in case of a failure during the resynchronization process. When the session is suspended, she can issue a Recover command to create a consistent version of the data on the H2 volumes. After the Recover command completes, she can issue the Start H1->H2 command to resynchronize the data from the San Francisco site to the Oakland site. If a failure occurs before her restarted session is in the Prepared state, she has a consistent version of the data on the H2 volumes. She must simply issue the Recover command to put the session into Target Available state and make the H2 volumes accessible from her servers. If the session was not in the Prepared state when it suspended, the subsequent Recover command does not issue the FlashCopy files to put the data on the H2 volumes. This means that the consistent data on the H2 volumes are not overwritten if the data to be copied to them is not consistent.

Metro Mirror Failover/Failback enabled for Open HyperSwap

Jane wants to run a Metro Mirror with Practice from San Francisco to Oakland. She wants to verify her recovery procedure for the Oakland site, but she cannot afford to stop running her Metro Mirror session while she takes time to practice a recovery. By using a Metro Mirror with Practice session, Jane is able to practice her disaster recovery scenario in Oakland while her Metro Mirror session runs uninterrupted. By practicing running her applications at the Oakland site, Jane is better prepared to make a recovery if a disaster ever strikes the San Francisco site.

While her session is running in a Prepared state, Jane practices a recovery at her Oakland site by issuing the Flash command. This command momentarily pauses the session and starts a FlashCopy to the H2 volumes. As soon as the FlashCopy is started, the session is restarted. These FlashCopy files create a consistent version of the data on the H2 volume that she can use for recovery testing, while her session continues to replicate data from San Francisco to Oakland. As a result, she can carry out her recovery testing without stopping her replication for any extended duration of time.

If the Metro Mirror session suspends due to a failure, Jane can use the practice session to restart her data replication while maintaining a consistent copy of the data at the Oakland site, in case of a failure during the resynchronization process. When the session is suspended, she can

issue a Recover command to create a consistent version of the data on the H2 volumes. After the Recover command completes, she can issue the Start H1->H2 command to resynchronize the data from the San Francisco site to the Oakland site. If a failure occurs before her restarted session is in the Prepared state, she has a consistent version of the data on the H2 volumes. She must simply issue the Recover command to put her session into Target Available state and make the H2 volumes accessible from her servers. If the session was not in the Prepared state when it suspended, the subsequent Recover command does not issue the FlashCopy files to put the data on the H2 volumes. This means that the consistent data on the H2 volumes are not overwritten if the data to be copied to the volumes is not consistent.

Selecting a HyperSwap session

A global insurance company has elected to use Tivoli Storage Productivity Center for Replication to manage its disaster recovery environment. Jane wants minimal data exposure, both for planned outages such as routine maintenance, and for unplanned disasters. They have CKD volumes on System Storage DS8000 devices, and use z/OS mainframes. They have two data centers in New York.

Jane reviews the Tivoli Storage Productivity Center for Replication documentation, and chooses a Metro Mirror recovery solution, since her priority is being protected against regional disasters. Jane decides to use Metro Mirror solution, because her company has two data centers located near each other. Jane realizes that because she uses z/OS, CKD, and System Storage DS8000 hardware, she is also able to use a HyperSwap solution. Using Metro Mirror Failover/Failback with HyperSwap, Jane can minimize application impact, while maintaining seamless failover to her secondary site. Jane decides Metro Mirror Failover/Failback with HyperSwap is best for the needs of her company.

After installing and configuring Tivoli Storage Productivity Center for Replication on z/OS, Jane starts the Tivoli Storage Productivity Center for Replication GUI. She adds the Tivoli Storage Productivity Center for Replication storage devices she intends to use on all three sites. From the Session Overview panel, Jane launches the Create Session wizard, and selects the **Metro Mirror Failover/Failback** session type. Then, she continues the wizard, she selects the **Manage H1H2 with HyperSwap** option. After finishing the wizard, Jane clicks **Launch Add Copy Sets Wizard**. She completes this wizard, and issues a **Start H1->H2** command. After the initial copy is completed, Jane is safely replicating her data between both sites. She can also issue a HyperSwap between sites 1 and 2, enabling her to switch sites with minimal application impact during either a disaster or maintenance period.

Performing a planned HyperSwap

Jane's company has successfully been using Metro Mirror Failover/Failback with HyperSwap sessions for the past three months. However, Jane needs to perform maintenance on an H1 box. During this time, Jane does not want her applications or replication to be interrupted. To prevent this from happening, shortly before the maintenance is scheduled to begin, Jane decides to use the Tivoli Storage Productivity Center for Replication GUI to perform a HyperSwap to the H2 volumes. This transitions the applications so that they write to H2. To perform a planned HyperSwap, Jane issues a HyperSwap command.

Understanding what happens when an unplanned HyperSwap occurs

Several weeks after the planned maintenance at Jane's company is completed, an incident occurs at the H1 site. A disk controller fails, causing one of the H1 volumes to encounter a permanent I/O error. Fortunately, Jane's data is safe because she used Metro Mirror Failover/Failback with HyperSwap, and her H2 volume is an exact duplicate of the H1 volume. When the permanent I/O error is detected, a HyperSwap is triggered. The application seamlessly transitions to writing to the H2 volumes. Her data is safe, and her applications are not interrupted.

Jane configured an Simple Network Management Protocol (SNMP) listener to alert her to any events, so she receives the SNMP event indicating that a HyperSwap occurred. Jane investigates the cause of the HyperSwap and uses the z/OS console to identify the volume that triggered the HyperSwap. Jane replaces the faulty disk controller. Then, to recover from the unplanned HyperSwap, Jane issues the **Start H2->H1** command.

Global Mirror


Global mirror is a method of asynchronous, remote data replication between two sites that are over 300 kilometers (km) apart. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system.

The data on the target is typically written a few seconds after the data is written to the source volumes. When a write is issued to the source copy, the change is propagated to the target copy, but subsequent changes are allowed to the source before the target verifies that it has received the change. Because consistent copies of data are formed on the secondary site at set intervals, data loss is determined by the amount of time since the last consistency group was formed. If your system stops, Global Mirror might lose some data that was being transmitted when the disaster occurred. Global Mirror still provides data consistency and data recoverability in the event of a disaster.

Global Mirror Single Direction

A Global Mirror single direction session allows you to run your Global Mirror replication from only the primary site.

For ESS, DS6000, and DS8000 storage systems, each copy set in the Global Mirror single direction session consists of two host volumes and a journal volume. The following figure illustrates how a Global Mirror single direction session works on an ESS, DS6000, and DS8000 storage systems:



For the following storage systems, each copy set in the Global Mirror Single Direction session consists of two host volumes:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

The following figure illustrates how a Global Mirror Single Direction session works on these storage systems:




Global Mirror Either Direction with Two-Site Practice (ESS, DS6000, and DS8000)

A Global Mirror Either Direction with Two Site Practice session allows you to run Global Mirror replication from either the primary or secondary site. It combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on a remote site at a distance over 300 km away from your first site. This practice session allows you to create practice volumes on both the primary and secondary site to practice what you might do if a disaster occurred, without losing your disaster recovery capability.

Note: This replication method is available on only ESS, DS6000, and DS8000 storage systems.


The session consists of two host volumes, two intermediate volumes, and two journal volumes. The following figure illustrates how a Global Mirror either direction with two-site practice session works:



Global Mirror Failover/Failback

Using Global Mirror failover/failback, your data exists on the second site that is more than 300 km away, and you can use failover/failback to switch the direction of the data flow. This enables you to run your business from the secondary site.

For ESS, DS6000, and DS8000 storage systems, each copy set in the Global Mirror failover/failback session consists of two host volumes and a journal volume. The following figure illustrates how a Global Mirror failover/failback session works on an ESS, DS6000, or DS8000 storage systems:



For the following storage systems, each copy set in the Global Mirror failover/failback session consists of two host volumes:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified
- The XIV system


The following figure illustrates how a Global Mirror failover/failback session works on these storage systems.



Global Mirror Failover/Failback with Practice

A Global Mirror failover/failback with practice session combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on a remote site at a distance over 300 km away from your primary site. You can use this to practice what you might do if a disaster occurred, without losing your disaster recovery capability.


For ESS, DS6000, and DS8000 storage systems, each copy set in the Global Mirror failover/failback with practice session consists of two host volumes, an intermediate volume, and a journal volume. The following figure illustrates how a Global Mirror failover/failback with practice session works on an ESS, DS6000, or DS8000 storage system:



For the following storage systems, each copy set in the Global Mirror failover/failback with Practice session consists of two host volumes and an intermediate volume:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

The following figure illustrates how a Global Mirror Failover/Failback with Practice session works on these storage systems:



Examples

Global Mirror Single Direction

Although Jane's FlashCopy and Metro Mirror copies were both planned, Jane realizes that sometimes unforeseen things happen, and she wants to make sure her data is safe. Because Jane works in San Francisco, she wants her other site to be far away in case of a localized disaster. Her other site is based in Houston. Jane's foresight pays off when a minor earthquake occurs in San Francisco and power and communications both go down. Fortunately, Jane has arranged for the data on customer accounts that have recently opened or closed to be asynchronously copied in Houston, using

Global Mirror. Jane risks losing the bytes of data that were being processed when the tremor disrupted the San Francisco process, but she views that as a minor inconvenience when weighed next to the value of backing up her data in a non-earthquake zone.

Global Mirror with Practice

Jane wants to run a Global Mirror with practice from San Francisco to Houston. She wants to verify her recovery procedure for the Houston site, but she cannot afford to stop running her Global Mirror session while she takes time to practice a recovery. By using a Global Mirror with Practice session, Jane is able to practice her disaster recovery scenario in Houston while her Global Mirror session runs uninterrupted. By practicing running her applications at the Houston site, Jane will be better prepared to make a recovery if a disaster ever strikes the San Francisco site.

Global Mirror Either Direction with Two-Site Practice

Jane wants to run a Global Mirror with practice from San Francisco to Houston. She wants to verify her recovery procedure for the Houston site, but she cannot afford to stop running her Global Mirror session while she takes time to practice a recovery. By using a Global Mirror either direction two-site with practice session, Jane is able to practice her disaster recovery scenario in Houston while her Global Mirror session runs uninterrupted. By practicing running her applications at the Houston site, Jane will be better prepared to make a recovery if a disaster ever strikes the San Francisco site.

Jane can use the Global Mirror either direction with two-site practice session to run asynchronous consistent data replication from either the San Francisco site or the Houston site. (She can practice her disaster recovery at the target site, no matter what her current production site is.) Jane's business is able to run a consistent Global Mirror session from its Houston site back to San Francisco while running a production at the Houston site.

Setting up Global Mirror for Resource Groups on System Storage DS8000

If resource groups are defined on a System Storage DS8000, Global Mirror session IDs might be defined for some users. Tivoli Storage Productivity Center for Replication does not automatically determine which session IDs are valid. To determine which session IDs are valid, you must modify the `rmserver.properties` file and add the following property:

```
gm.master.sessionid.gm_role,session_name = xx
```

where `gm_role` is the role that has the master volume (for example, H1 in a Global Mirror failover/failback session), `session_name` is the name of the session that uses the session ID, and `xx` is the decimal number for the session ID.

Important: System Storage DS8000 represents session IDs as a two-digit hexadecimal number. Use the decimal version of that number. For example, if you want a Global Mirror failover/failback session to use a session ID of 0F, the decimal number is 15 as shown in the following example:

```
gm.master.sessionid.H2.11194_wprac=15
```

Metro Global Mirror (ESS 800 and DS8000)

Metro Global Mirror is a method of continuous, remote data replication that operates between three sites that varying distances apart. Metro Global Mirror

combines Metro Mirror synchronous copy and Global Mirror asynchronous copy into a single session, where the Metro Mirror target is the Global Mirror source.

Notes:

- This replication method is available on only ESS 800 and DS8000 storage systems.
- You can select ESS 800 storage systems in only the H1 volume role. All other volume roles must use DS8000 volumes.
- You can mix ESS 800 and DS8000 volumes in the H1 volume role. If ESS 800 and DS8000 storage systems are both used in the H1 role, the DS8000 storage system performs Incremental Resync (IR), and the ESS 800 storage system performs a full copy. Because ESS 800 does not support the IR function, a full copy is required when changing from H1->H2->H3 to H1>H3 and from H2->H1->H3 to H2->H3.

Metro Global Mirror maintains a consistent copy of data at the remote site, with minimal impact to applications at the local site. This remote mirroring function works in combination with FlashCopy to meet the requirements of a disaster-recovery solution by providing the following features:

- Fast failover and failback
- Rapid reestablishment of three-site mirroring, without production outages
- Data currency at the remote site with minimal lag behind at the local site, an average of only 3 - 5 seconds for many environments
- Quick resynchronization of mirrored sites using only incremental changes

If IBM Tivoli Storage Productivity Center for Replication is running on z/OS, you can configure a Metro Global Mirror session to control the Metro Mirror relationship between the primary and secondary site using HyperSwap. With HyperSwap enabled, a failure on the primary storage system causes an automatic HyperSwap, transparently redirecting application I/O to the auxiliary storage system. The Global Mirror relationship continues to run uninterrupted throughout this process. With this configuration, you can achieve near zero data loss at larger distances.

Using synchronous mirroring, you can switch from local site H1 to remote site H2 during a planned or unplanned outage. It also provides continuous disaster recovery protection of site H2 through site H3, without the necessity of additional reconfiguration, if a switch from site H1 occurs. With this configuration, you can reestablish H2->H1->H3 recoverability while production continues to run at site H2. Additionally, this cascaded setup can reduce the load on-site H1 as compared to some multi-target (non-cascaded) three-site mirroring environments.


Important:

- If HyperSwap occurs by event when running a Metro Global Mirror with a HyperSwap session, a full copy of the data occurs to return to a full three-site configuration. If you issue a HyperSwap command when running a Metro Global Mirror with a HyperSwap session, a full copy does not occur. A full copy is required only for an unplanned HyperSwap or a HyperSwap initiated using the z/OS SETHS SWAP command.
- In Metro Global Mirror and Metro Global Mirror with Practice sessions, when the H1 is on an ESS 800, you might risk filling up the space efficient journal volumes. Because incremental resynchronization is not supported on the ESS 800, full copies are performed in many of the transitions.

Metro Global Mirror

A Metro Global Mirror session with Practice combines Metro Mirror, Global Mirror, and FlashCopy across three sites to provide a point-in-time copy of the data on the third site.

The following figure illustrates how a Metro Global Mirror session works.




Metro Global Mirror with Practice

A Metro Global Mirror session with Practice combines Metro Mirror, Global Mirror, and FlashCopy across three sites to provide a point-in-time copy of the data on the third site. You can use this session to practice what you might do if a disaster occurred without losing your disaster recovery capability.

The session consists of three host volumes, an intermediate volume, and a journal volume. The following figure illustrates how a Metro Global Mirror with Practice session works.

Note: A Metro Global Mirror with practice session can be created when the three-site license has been applied to the server.



Note: In Metro Global Mirror and Metro Global Mirror with Practice sessions, when the H1 is on an ESS 800, you might risk filling up the space efficient journal volumes. Because incremental resynchronization is not supported on the ESS 800, full copies are performed in many of the transitions.

Examples

Metro Global Mirror

Although Jane works in San Francisco, she wants to give herself the ability to run her business from either Oakland (her secondary site) or Houston (her tertiary site). Jane can use Metro Global Mirror with failover/failback to switch the direction of the data flow, so that she can run her business from either Oakland or Houston. Metro Global Mirror means that Jane has zero data loss backup at her secondary site, and minimal data loss at her tertiary site.

Metro Global Mirror with Practice

Jane wants to run a Metro Global Mirror with Practice from San Francisco to Houston. She wants to verify her recovery procedure for the Houston site. However, she cannot afford to stop running her Metro Global Mirror session while she takes time to practice a recovery. By using a Metro Global Mirror with Practice session, Jane is able to practice her disaster recovery scenario in Houston while her Metro Global Mirror session runs uninterrupted. By practicing running her applications at the Houston site, and being prepared to run her applications at the Oakland site if necessary, Jane will be better prepared to make a recovery if a disaster ever strikes the San Francisco site.

Jane can use Metro Global Mirror with Practice to switch the direction of the data flow, so that she can run her business from either Oakland or Houston. Using Metro Global Mirror, Jane has zero data loss backup at her secondary site; and minimal data loss at her tertiary site.

Managing a session with HyperSwap and Open HyperSwap replication

HyperSwap and Open HyperSwap provide high availability of data in the case of a primary disk storage system failure. When a failure occurs in writing input/output (I/O) to the primary storage system, the failure is detected by IOS, and IOS automatically swaps the I/O to the secondary site with no user interaction and little or no application impact.

Sessions that can be enabled for HyperSwap or Open HyperSwap:

You can create sessions that have HyperSwap or Open HyperSwap capabilities. Enabling swapping provides a session with a highly available business continuity solution.

Sessions that can enable HyperSwap

The following session types can enable HyperSwap:

- Basic HyperSwap
- Metro Mirror with Failover/Failback
- Metro Global Mirror

To enable HyperSwap, the following circumstances must apply:

- The session is running on an Tivoli Storage Productivity Center for Replication server that is running on IBM z/OS.

- The volumes are only for TotalStorage Enterprise Storage Server, System Storage DS8000, and DS6000 systems.
- The volumes are count key data (CKD) volumes that are attached to the z/OS system.

Sessions that can enable Open HyperSwap

Only Metro Mirror with Failover/Failback session type can enable Open HyperSwap.

To enable Open HyperSwap, the following circumstances must apply:

- The volumes in the session are only for System Storage DS8000 5.1 or later volumes.
- The volumes in the session are fixed block and mounted to IBM AIX 5.3 or AIX 6.1 hosts with the following modules installed:
 - Subsystem Device Driver Path Control Module (SDDPCM) version 3.0.0.0 or later
 - Multi-Path Input/Output (MPIO) module (the version that is provided with AIX version 5.3 or 6.1)
- The connections between the AIX host systems and the Tivoli Storage Productivity Center for Replication server have been established.

Setting up the environment for HyperSwap:

You must set up an environment that supports HyperSwap before attempting to enable HyperSwap for a IBM Tivoli Storage Productivity Center for Replication session.

The following steps must be completed before HyperSwap can be enabled. For more information about these steps, see the IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide

1. Install IBM Tivoli Storage Productivity Center for Replication for System z.
2. Perform the post installation tasks of setting up the data store and other necessary system settings.
3. Ensure that all RESERVEs are converted to global enqueues (ENQs).
4. Ensure that all volumes in the session that you are enabling for HyperSwap are attached to the IBM z/OS system that is running Tivoli Storage Productivity Center for Replication.

Setting up the environment for Open HyperSwap:

You must set up an environment that supports Open HyperSwap before attempting to enable Open HyperSwap for a IBM Tivoli Storage Productivity Center for Replication session.

The following steps must be completed before Open HyperSwap can be enabled:

1. Ensure that the IBM AIX hosts and IBM System Storage DS8000 to meet the following hardware and software requirements:

AIX requirements

Open HyperSwap support requires AIX version 5.3 or 6.1. (You can find the supported AIX version for each Tivoli Storage Productivity Center for Replication release in the support matrix at <http://www-01.ibm.com/support/docview.wss?rs=40&context=SSBSEX&context=SSMN28&context=SSMMUP&context=SS8JB5>)

&context=SS8JFM&uid=swg21386446&loc=en_US&cs=utf-8&lang=en.
Click the link for the applicable release under Agents, Servers and GUI.)

You must have the following AIX modules installed:

- Subsystem Device Driver Path Control Module (SDDPCM) version 3.0.0.0 or later
- Multi-Path Input/Output (MPIO) module (the version that is provided with AIX version 5.3 or 6.1)

System Storage DS8000 hardware requirements

Only System Storage DS8000 storage systems are supported. Open HyperSwap requires System Storage DS8000 5.1 or later.

Open HyperSwap does not support High Availability Cluster Multi-Processing (HACMP[™]).

2. Create connections from Tivoli Storage Productivity Center for Replication to the AIX hosts (see “Adding a host system connection” on page 666).
3. Assign copy set volumes from the storage device to the host using the System Storage DS8000 command-line interface (CLI) or graphical user interface (GUI).
4. Run the AIX `cfgmgr` command to discover the volumes assigned to the host.

Considerations for Open HyperSwap and the AIX host:

- A single session that has Open HyperSwap enabled can manage multiple hosts; however, each host can be associated with only one session. Multiple hosts can share the same session.
- For AIX 5.3, a single host can manage a maximum of 1024 devices that have been enabled for Open HyperSwap on the host, with 8 logical paths configured for each copy set in the session. For AIX 6.1, a single host can manage a maximum of 1024 devices that have been enabled for Open HyperSwap on the host, with 16 logical paths configured for each copy set in the session.
- If an application on the host has opened a device, a Tivoli Storage Productivity Center for Replication session for that device cannot be terminated. The **Terminate** command fails. To terminate the session, you must either close the application or remove the copy sets from the session. If you remove copy sets from the session, you must ensure that the application writes to the correct volume when the copy set relationship is restored.
- It is possible for Open HyperSwap to fail on a subset of hosts for the session and work on the remaining hosts for the same session. In this situation, you must determine the best action to take if the application is writing to volumes on the source system as well as volumes on the target system. Contact the IBM Support Center if you need assistance determining the best solution for this issue.
- To enable support for Open HyperSwap on the host, refer to the *IBM System Storage Multipath Subsystem Device Driver User's Guide*.

Configuring timers to support Open HyperSwap:

There are configurable timeout values for the storage system, IBM Tivoli Storage Productivity Center for Replication, and IBM AIX hosts systems that can affect the operation of Open HyperSwap.

The following list describes the various timeout values that can affect Open HyperSwap:

Storage system quiesce timeout value

This quiesce timeout timer begins when the storage system starts a quiesce operation. When the timer value expires, input/output (I/O) is resumed on the primary device. The default timeout value is two minutes, but the value can be set from 30 to 600 seconds. To set the quiesce timeout value, see the information about the **chdev** command in the *IBM System Storage Multipath Subsystem Device Driver User's Guide*.

Storage system long busy timeout value

This timeout value is the time in seconds that the logical subsystem (LSS) consistency group volume stays in the long busy state after a remote mirror and copy error is reported.

Timeout values for the applications that are on the host

The various applications that are running on the host have timeout values. The timeout values vary depending on the application.

Considerations for setting timers

Consider the following information for setting timers:

- If the host quiesce timer is set to a shorter value than the Tivoli Storage Productivity Center for Replication response timer, it is possible that an I/O swap failure can occur. If a storage system triggers an unplanned failover and if the storage system quiesce timer expires before Tivoli Storage Productivity Center for Replication responds, the host attempts to write I/O to the primary volume where the loss of access occurred. If the hardware condition that caused the loss of access continues, the attempt to write I/O fails again and an unplanned Open HyperSwap is not performed.
- If the host quiesce timer is set to a longer value than the Tivoli Storage Productivity Center for Replication response timer, an application timeout might occur if Open HyperSwap takes too long to complete.

Enabling a Session for HyperSwap or Open HyperSwap:

Enabling HyperSwap or Open HyperSwap for a session provides a combined business recovery and business continuity solution.

To ensure that your environment supports HyperSwap or Open HyperSwap, see “Setting up the environment for HyperSwap” on page 101 or “Setting up the environment for Open HyperSwap” on page 101.

Perform these steps to enable HyperSwap or Open HyperSwap for a session.

1. In the IBM Tivoli Storage Productivity Center for Replication navigation tree, select **Sessions**. Click the radio button next to the session that you want to enable.
2. From the **Select Action** menu, select **View/Modify Properties** and click **Go**. If you have not created the session, click **Create Session**. You can enable HyperSwap or Open HyperSwap on the Properties page
3. Under **ESS / DS Metro Mirror Options**, select from the following HyperSwap or Open HyperSwap options:
 - **Manage H1-H2 with HyperSwap**. This option enables a session to manage the H1-H2 sequence using HyperSwap. If you select this option, select from the following additional options.
 - **Disable HyperSwap**. Select this option to prevent a HyperSwap from occurring by command or event.

- **On Configuration Error.** Choose one of the following options:
 - **Partition the system(s) out of the sysplex.** Select this option to partition out of the sysplex when a new system is added to the sysplex and encounters an error in loading the configuration. A restart of your system is required if you select this option.
 - **Disable HyperSwap.** Select this option to prevent a HyperSwap from occurring by command or event.
 - **On Planned HyperSwap Error.** Choose one of the following options:
 - **Partition out the failing system(s) and continue swap processing on the remaining system(s).** Select this option to partition out the failing system and continues the swap processing on any remaining systems.
 - **Disable HyperSwap after attempting backout.** Select this option to stop the HyperSwap action, and disable the HyperSwap commands or events.
 - **On Unplanned HyperSwap.** Choose one of the following options:
 - **Partition out the failing system(s) and continue swap processing on the remaining system(s).** Select this option to partition out the failing systems and continues the HyperSwap processing on the remaining systems when a new system is added to the sysplex and HyperSwap does not complete. A restart of your system is required if you select this option.
 - **Disable HyperSwap after attempting backout.** Select this option to stop the HyperSwap action, and disable the HyperSwap commands or events.
 - **Manage H1-H2 with Open HyperSwap.** If volumes are attached to an IBM AIX host, Tivoli Storage Productivity Center for Replication can manage the H1-H2 sequence of a Metro Mirror session using Open HyperSwap. If this option is selected, a failure on the host accessible volumes triggers a swap, which redirects application I/O to the secondary volumes. Only volumes that are currently attached to the host systems that are defined on the Tivoli Storage Productivity Center for Replication Host Systems panel are eligible for Open HyperSwap.
 - **Disable Open HyperSwap.** Select this option to prevent a swap from occurring by a command or event while keeping the configuration on the host system and all primary and secondary volumes coupled.
4. Click **OK** to apply the selected options.

Restarting an AIX Host System that is enabled for Open HyperSwap:

When an IBM AIX host system is restarted, the host automatically attempts to open any volumes for input/output (I/O) that were open prior to the restart. If Open HyperSwap was enabled for a set of volumes on the host system, the host must determine which storage system is the primary system before the host can allow the volumes to be opened.

If the Metro Mirror relationship for the set of volumes is in a Prepared or Suspended state and the host has connectivity to both the primary and secondary storage systems, the host can determine through the hardware which storage system is the primary system. In this situation, the host automatically opens the volumes.

If the Metro Mirror relationship for the set of volumes is in a Prepared state and the host has connectivity to only the secondary storage system, all I/O to the

volumes might be blocked on the host system until the host is able to verify the primary volume in the relationship. The AIX command varyonvg will fail to open the volumes for I/O to prevent the application from writing to the incorrect site. If the host can determine which volume is the primary volume in the relationship and connectivity to the primary storage system is still lost, a Hyperswap event is triggered. This event causes all I/O to be automatically opened and directed to the secondary storage system.

If the Metro Mirror relationship for the set of volumes is in a Target Available state after a Hyperswap or a Recover command has been issued for the session, or the host system does not have the connectivity necessary to determine which site is the primary site and all I/O to the volumes are blocked on the host system. The varyonvg command will fail to open the volumes for I/O to prevent the application from writing to the incorrect site.

Unblocking I/O on the host system after a host system restart

When any of the previous scenarios cause I/O to be blocked, manual actions might be necessary to remove the block.

If the relationships are in a Target Available state on the hardware, issue a Start command to the session in the desired direction of the relationship. This action defines the primary storage system for the host. The host system can allow the volumes to be opened to I/O.

If the relationships cannot be restarted, or the host cannot determine the primary storage system, it might be necessary to manually decouple the volumes on the host system.

To decouple the volumes the following options are available:

- Option 1: Terminate the session or remove the copy set. This option would require a full copy when the relationships are restarted.
- Option 2: Remove Object Data Manager (ODM) entries using the following command:

```
odmdelete -o Volume_Equivalency
```

CAUTION:

This command should be used only for this scenario because the command deletes copy set information.

Planned and unplanned swaps:

Once a session has been enabled for HyperSwap or Open HyperSwap and reaches the Prepared state, IBM Tivoli Storage Productivity Center for Replication loads the configuration of volumes that are capable of being swapped onto IBM z/OS or AIX.

When the load is complete, the session is capable of a planned or unplanned swap. The H1-H2 role pair on the session show a type of HS. An H is displayed over the connection in the dynamic image for that role pair as shown in the following

Metro Mirror Failover/Failback




figure.

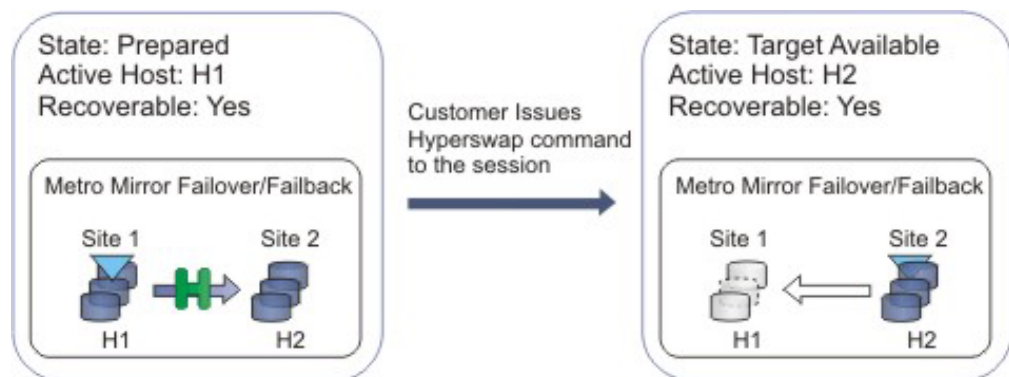
Performing a Planned Swap

Once the session configuration is loaded on z/OS for HyperSwap or AIX for Open HyperSwap, the session is considered swap capable. There may be cases such as a planned maintenance or a migration from the primary storage, in which a planned swap might be required. Once the session is in a swap capable state, a planned swap can be executed by issuing the HyperSwap command against the session.

Once a planned swap is run for z/OS HyperSwap and Open HyperSwap, the session is transitioned to a Target Available state and all the H1-H2 pairs are in a Target Available state. If the H1-H2 role pair was consistent at the time of the swap, the session will have a status of Normal and will indicate that H1-H2 is consistent. If the H1-H2 role pair was not consistent at the time of the swap, the session might display a status of SEVERE because the session is inconsistent. The active host on the session is then displayed as H2.

All input/output (I/O) should have been redirected to the H2 volumes. After a successful swap to site 2, it is not possible to re-enable copy to site 2. Therefore, it is not possible to issue a Start H1->H2 command. The only way to restart the copy is a Start H2->H1 command. To have the volumes protected with high availability and disaster recovery again, the error that caused the swap must be fixed and then the session must be manually restarted to begin copying to the other site.

The following figure illustrates a planned swap.



What happens when an unplanned swap occurs


Once the session configuration is loaded on z/OS for HyperSwap or AIX for Open HyperSwap, the session is considered swap capable. In the event of a primary I/O error, a swap occurs automatically. For HyperSwap, z/OS performs the entire swap

and then alerts Tivoli Storage Productivity Center for Replication that a swap has occurred. For Open HyperSwap, Tivoli Storage Productivity Center for Replication and the AIX host work together to perform the swap.

Once an unplanned swap occurs for HyperSwap and Open HyperSwap, the session is transitioned to a Target Available state and all the H1-H2 pairs are in a Target Available state. If the H1-H2 role pair was consistent at the time of the swap, the session will have a Status of Normal and will indicate that H1-H2 is consistent. If the H1-H2 role pair was not consistent at the time of the swap, the session might display a status of SEVERE because the session is inconsistent. The active host on the session is then displayed as H2.

All I/O should have been redirected to the H2 volumes. After a successful swap to site 2, it is not possible to re-enable copy to site 2. Therefore, it is not possible to issue a Start H1->H2 command. The only way to restart the copy is a Start H2->H1 command. To have the volumes protected with high availability and disaster recovery again, the error that caused the swap must be fixed and then the session must be manually restarted to begin copying to the other site.

The following figure illustrates an unplanned swap.



Scenarios requiring a full copy in Metro Global Mirror with HyperSwap sessions

In the following cases, a full copy is required to return to the three-site configuration after a swap:

- If you are running a Metro Global Mirror session with HyperSwap and you issue the HyperSwap command using the z/OS HyperSwap API rather than the Tivoli Storage Productivity Center for Replication graphical user interface (GUI).
- If you are running a Metro Global Mirror session with HyperSwap and an unplanned swap occurs.

Verifying that a session is capable of a planned or unplanned swap:

You can verify whether a sessions is capable of a planned or unplanned swap from the IBM z/OS console (HyperSwap) or the IBM AIX host (Open HyperSwap).

Perform these steps to verify the status of HyperSwap from the z/OS console:

1. Issue the `ds hs,status` command for the overall status of the HyperSwap session. For example:

```

15.03.06 SYSTEM1 d hs,status
15.03.06 SYSTEM1 STC00063 IOSHM0303I HyperSwap Status 531
Replication Session: SR_HS
HyperSwap enabled
New member configuration load failed: Disable
Planned swap recovery: Disable
Unplanned swap recovery: Disable
FreezeAll: No
Stop: No

```

2. Issue the `ds hs,config(detail,all)` command to verify all the volumes in the configuration. For example:

```

15.03.51 SYSTEM1 d hs,config(detail,all)
15.03.51 SYSTEM1 STC00063 IOSHM0304I HyperSwap Configuration 534
Replication Session: SR_HS
Prim. SSID UA DEV# VOLSER Sec. SSID UA DEV# Status
06 02 00F42 8K3602 06 04 00FA2
06 01 00F41 8K3601 06 03 00FA1
06 00 00F40 8K3600 06 02 00FA0

```

Perform these steps to verify the status of Open HyperSwap from the AIX host:

1. Issue the `pcmpath query device` command to see the session association and which path the input/output (I/O) is currently being routed to, which is indicated by an asterisk. For example:

```

host1> pcmpath query device 14

DEV#: 14 DEVICE NAME: hdisk14 TYPE: 2107900 ALGORITHM: Load Balance
SESSION NAME: session1
OS Direction: H1<-H2
=====
PRIMARY SERIAL: 25252520000
-----
Path# Adapter/Path Name State Mode Select Errors
0 fscsi0/path0 CLOSE NORMAL 6091 0
1 fscsi0/path2 CLOSE NORMAL 6300 0
2 fscsi1/path4 CLOSE NORMAL 6294 0
3 fscsi1/path5 CLOSE NORMAL 6187 0

SECONDARY SERIAL: 34343430000 *
-----
Path# Adapter/Path Name State Mode Select Errors
4 fscsi0/path1 CLOSE NORMAL 59463 0
5 fscsi0/path3 CLOSE NORMAL 59250 0
6 fscsi1/path6 CLOSE NORMAL 59258 0
7 fscsi1/path7 CLOSE NORMAL 59364 0

```

Temporarily disabling HyperSwap or Open HyperSwap:

In some situations, it might be necessary to temporarily disable the HyperSwap or Open HyperSwap capabilities for a session.

You might want to disable HyperSwap or Open HyperSwap under the following circumstances:

- Performing maintenance
- Inability for one sysplex member to communicate with one or more volumes

Perform these steps to disable Open HyperSwap for a specific session:

1. In the navigation IBM Tivoli Storage Productivity Center for Replication tree, select **Sessions**. The Sessions panel is displayed.
2. Select the sessions for which you want to disable HyperSwap or Open HyperSwap.

3. Select **View/Modify Properties** from the **Select Actions** list, and click **Go**.
4. Select **Disable HyperSwap** or **Disable Open HyperSwap** and click **OK**.

Tip: On management servers that run IBM z/OS, you can also disable HyperSwap from an MVS command prompt by entering **SETHS DISABLE**.

Using active and standby Tivoli Storage Productivity Center for Replication servers with HyperSwap or Open HyperSwap:

To ensure that there is an IBM Tivoli Storage Productivity Center for Replication server available in the event of a disaster, an active and standby management server configuration can be set up. You can enable HyperSwap and Open HyperSwap for sessions while maintaining an active and standby server configuration.

Active and standby servers with HyperSwap

When the storage system is set up to connect through the z/OS interface, the connection information is automatically sent to the standby server and a connection is attempted. It is possible that the connection will fail if the standby server is not running on z/OS or does not have access to the same volumes. If the connection fails, any takeover done on the standby server will not be able to manage the HyperSwap. On z/OS, if the session configuration had been successfully loaded prior to the HyperSwap, the z/OS system is still capable of performing the HyperSwap. If the z/OS system swaps the volumes but there is no communication possible to the Tivoli Storage Productivity Center for Replication server, the session recognizes that the pairs became suspended and the session will go into a Suspended/Severe state. From this state the customer can clear the **Manage H1-H2 with Hyperswap** option and issue the Recover command for the session to get the session into a Target Available state.

Active and standby servers with Open HyperSwap

When there is an active and standby management server configuration and a host system connection is added to the active server, the host system connection is automatically sent to the standby server and a connection is attempted. Once the session configuration is loaded on IBM AIX, Open HyperSwap is possible only if there is continual communication between AIX and the Tivoli Storage Productivity Center for Replication server. If a takeover is performed on a standby server that is unable to connect to the host system that is managing the swap, the session is no longer Open HyperSwap capable. Communication to the host system must be activated before the session can become Open HyperSwap capable again.

Related tasks:

“Managing management servers” on page 643

This section provides information about how to set up active and standby management servers, restore a lost connection between the management servers, or perform a takeover on the standby management server.

Alerts

You can set up Tivoli Storage Productivity Center so that it examines the data that it collects about your storage infrastructure and writes an alert to a log when an event occurs. You also can specify that an action is initiated, such as sending an SNMP trap, sending an email, or running a script when the event occurs.

Alerts that are triggered based on the data collected by monitoring jobs (for example, pings, scans, and probes) must be defined before the monitoring jobs are run. Alerts that are triggered based on an external source, such as an SNMP trap, must be defined after the SNMP server is configured properly.

You can define an alert in the following ways:

- As part of a data collection schedule
- As part of an alert definition

When an event occurs and triggers an alert, the alert is written to a log. You can also select one or more other ways to be notified of the event. These alert notifications include SNMP traps, IBM Tivoli Enterprise Console events, Tivoli Storage Productivity Center login notifications, operating-system event logs, or email.

Note: Alerts are not generated in a Tivoli Storage Productivity Center instance for actions that you perform from that instance. For example, if you start Tivoli Storage Productivity Center and use Disk Manager to assign or unassign volumes for a subsystem, you do not receive alerts for those volume changes in that instance. However, if you assign and unassign volumes outside of that Tivoli Storage Productivity Center instance, an alert is generated.

Triggering conditions for alerts

In general, the following types of conditions can trigger alerts:

- A data collection schedule failed to complete.
- A change occurred in the storage infrastructure.
- A performance threshold was violated.

The specific conditions that can trigger events vary, depending on the type of storage resource that is being monitored. See “Triggering conditions for alerts” on page 1158 for more information about triggering conditions.

Alert notifications and triggered actions

Notification methods, or triggered actions, define the method by which you are notified of an alert: SNMP Trap, Tivoli Enterprise Console Event, Login Notification, Windows Event Log, UNIX Syslog, and email. Alerts are always written to the error log. Additionally, if a Data agent or Storage Resource agent is deployed on a monitored computer, you can run a script or start an IBM Tivoli Storage Manager job in response to the alert. See “Triggered actions for alerts” on page 1172 for more information about triggered actions.

Prerequisites for successfully using alerts

The following conditions must be met in order to successfully use alerts:

- Data collection schedules are configured and scheduled to run on a regular basis.
- If you want to be notified about an alert in some way other than an entry in the log file, such as using SNMP traps, Tivoli Enterprise Console events, or email, alert notification instructions must be configured before using the alert. See “Configuring alert destinations” on page 394).
- If an alert is triggered based on an SNMP trap from the monitored resource, you must properly configure the SNMP server of the monitor resource to enable

Tivoli Storage Productivity Center to listen to SNMP traps. The default port number is 162, and the default community is public.

Default alerts

The following types of alerts are set for your system by default:

- New entity found
- Job failed

You can change the triggered actions and enable or disable these alerts, but you cannot delete them.

Provisioning planner

Use the IBM Tivoli Storage Productivity Center *host planner*, the *storage planner*, and the *security planner* to provide guidance when you provision your storage system. The collective term for all three individual planners working together is the *provisioning planner*.

For example, to add more storage to a host, the three planners work together to create the storage, zone the new storage to the host, and configure the multipath settings on the host. After using the planners to perform these tasks, you can continue to perform typical Tivoli Storage Productivity Center jobs with the storage on the host, such as creating the file system.

You can call each planner individually. For example, you might modify an existing multipath or zone setting or create volumes on the subsystem for later use with the host planner.

Use the planners to plan an integrated set of high-level policies for performance, availability, and security, based on the current resource utilization. The set of selected policies determines the best action responses to performance, availability, and security events and constitutes the output plan, which is the record of the output of the planners. You can either accept plan advice to trigger an associated action (like creating a volume or performing zoning), or reject the advice and change the plan.

Modes of use

You can use the planners in the following ways:

Planning mode (Host planning advice)

Planning mode uses each of the planners: storage planner, provisioning planner, and security planner, to configure an end-to-end plan for storage connectivity. This results in a recommendation for zoning, volume assignment, and multipath software configuration.

Assist mode

Assist mode provides zoning recommendation that meet the connectivity requirements, as established in current policies.

Validate mode

Validate mode evaluates the current configuration against the connectivity requirements, using the configuration checker utility.

Host planner

Use the IBM Tivoli Storage Productivity Center *host planner* to provide guidance when you configure host to subsystem connectivity. The host planner accesses

default requirements or requirements that you set, and helps identify configurations that do not meet those requirements.

You can perform common network and storage provisioning tasks all at once by using the host planner with the other two planners, which are the *security planner* and the *storage planner*. The collective term for the three individual planners working together is *provisioning planner*.

For example, to add more storage to a host, the three planners work together to create the storage, zone the new storage to the host, and configure the multipath settings on the host. After using the planners to perform these tasks, you can continue to perform typical Tivoli Storage Productivity Center jobs with the storage on the host, such as creating the file system.

Host planning is the task of determining the connectivity between a host and the storage subsystems to obtain pathing and performance requirements. Inputs to the planning process are:

- The host
- Host ports and storage controller ports (and port performance utilization)
- Existing multipathing and zoning information
- The subsystem and identifiers
- The application profile (I/O, performance, and availability requirements)
- Pathing requirements (whether to use multipathing if it is available)

The planning process uses storage planning, configuration checking, and security planning for portions of the overall process.

Collection of information

The planner retrieves information from the Tivoli Storage Productivity Center database as its default input. You can provide explicit values to override default input. If you explicitly specify the number of paths, for example, the host planner verifies whether the value satisfies performance requirements. If not, the host planner calculates the number of paths, and which storage controller and host ports satisfy the requirements.

The fabric, disk and data discovery and probe operations must run to collect the information needed for the planner. If the information is not collected, the planner output identifies the limitation and presents a recommendation based on the subset of information that is available.

Required components

The host planner also uses the following components:

Data agent or Storage Resource Agent

The Tivoli Storage Productivity Center data agent is enhanced to deal with an SD multipathing driver on the host.

The enhanced data agent also collects SDD information, including SDD present, version, HBAs and ports, HBA info, vdisk and vpath configuration, and multipathing settings per vpath.

Multipathing control action application programming interfaces (APIs)

Unlike the storage planner or the security planner, control action APIs for multipathing do not exist.

Overview of the planning process

Host planning requires the following steps:

1. Validating that the host and subsystem are compatible by type and level of the host, the operating system, the subsystem, and the SDD software, for example
2. Determining the number of paths needed for the pathing policy and the expected I/O performance
3. Determining the connectivity required between the host and the subsystem
4. Determining the configuration for multipathing software
5. Configuring the zoning and SDD multipathing configuration

Policies related to the host planner

The host planner uses the following policies that you configure with the configuration checker utility:

Multipath

Specifies the multipath policy requirement (multiple or single path)

Preferred number of paths

Specifies required number of paths between host and subsystems

Multipath driver mode

Specifies multipath driver mode (fail-over, load balancing, round robin)

Redundant fabric

Specifies routing through isolated fabrics (must have at least 2 routes in separate fabrics, only 1 route required)

Planning guidance outputs

The host planner provides guidance that is based on the pathing and performance requirement inputs that it accesses.

Outputs from the planning process are:

- Validation that a plan is possible
- Zoning
- Subsystem assignment
- Multipathing (SDD) configuration information
- Configuration of zoning and SDD

Indication of violations

The planner notifies you of current policy violations before running a new plan that compensates for the violations. Violations are logged in the job log.

Command-line interface

There is no command-line interface for invoking the planning wizard.

Storage planner

Use the IBM Tivoli Storage Productivity Center *storage planner* to perform storage administration tasks, such as creating storage.

You can perform common network and storage provisioning tasks all at once by using the storage planner with the other two planners, which are the *host planner* and the *security planner*. The collective term for the three individual planners working together is *provisioning planner*.

For example, to add more storage to a host, the three planners work together to create the storage, zone the new storage to the host, and configure the multipath settings on the host. After using the planners to perform these tasks, you can continue to perform typical Tivoli Storage Productivity Center jobs with the storage on the host, such as creating the file system.

Collection of information

The storage planner collects and evaluates capacity, storage controller preferences, storage controller types, the number of volumes and volume size, performance requirements, the use of unassigned volumes, the RAID level and performance utilization and capacity utilization input from the Tivoli Storage Productivity Center database and from explicit values that you define.

Planning guidance outputs

The storage planner generates a plan that presents optimal storage controller and storage pool configurations.

Control action application programming interfaces (APIs)

The control action APIs are not implemented by the storage planner but are implemented by the device manager (DM).

Overview of the planning process

If you explicitly specify the storage pool and controller information, the storage planner checks to see whether the input performance and capacity requirements can be satisfied.

Supported devices

The storage planner provides support for IBM System Storage DS8000, IBM System Storage DS6000, and IBM Enterprise Storage (ES800 controllers). The storage planner is capable of working with any combination of these specific controllers.

Security planner

Use the IBM Tivoli Storage Productivity Center *security planner* to perform security administration tasks.

You can perform common network and storage provisioning tasks all at once by using the security planner with the other two planners, which are the *host planner* and the *storage planner*. The collective term for the three individual planners working together is *provisioning planner*.

For example, to add more storage to a host, the three planners work together to create the storage, zone the new storage to the host, and configure the multipath settings on the host. After using the planners to perform these tasks, you can continue to perform typical Tivoli Storage Productivity Center jobs with the storage on the host, such as creating the file system.

Collection of information

The security planner collects and evaluates existing host port information, storage controller ports, zones, zone sets, and switches from the Tivoli Storage Productivity Center database and from explicit values that you define as zoning input and volume masking input.

Planning guidance outputs

The security planner generates a plan that presents optimal zoning and volume masking.

Control action application programming interfaces (APIs)

The security planner control APIs are currently already implemented by the fabric and DM APIs.

Overview of the planning process

If you specify exact zoning and volume masking information details, the planner informs you whether your input is different than the existing configuration, or different than the optimal configuration that the planner calculates based on policies and the existing security parameters. If the optimal configuration differs from your selections, the planner proposes zoning and volume masking settings.

Reporting

This section contains conceptual information about reporting.

Use the Reporting facilities in the different managers of IBM Tivoli Storage Productivity Center to view both overview and detailed information about your storage. The over 500 available reports are constructed from the statistical information accumulated in the repository and enable you to select exactly the computers, storage systems, switches, clusters, disks, file systems, hypervisors, users and even other Tivoli Storage Productivity Center servers upon which you want to report. You can view reports as tables of data or as graphical charts, which provides you with the flexibility to get a very detailed look at your information or a very high level overview.

You must collect information about your environment before you can use reports to view details about the storage resources in it. You can use monitoring jobs in Tivoli Storage Productivity Center such as probes, scans, pings, and Tivoli Storage Productivity Center server probes, and discovery jobs to gather comprehensive information and statistics about your storage resources.

Fabric and zones

Use the IBM Tivoli Storage Productivity Center to learn more about fabrics and zones.

Zones, zone aliases, and zone sets

Zones, zone aliases, and zone sets allow logical grouping of ports and storage devices within a storage area network. This section describes zoning concepts and elements.

In a storage area network a *zone* is a logical grouping of ports to form a virtual private storage network. Zones that belong to a single SAN can be grouped into a

zone set, which can be activated or deactivated as a single entity across all switches in the fabric. A zone set can contain one or more zones, and a zone can be a member of more than one zone set. Using zoning, you can automatically or dynamically arrange fabric-connected devices into logical groups across a physical fabric.

Ports and devices in a zone are called *zone members*. A zone can contain one or more zone members. Ports that are members of a zone can communicate with each other, but they are isolated from ports in other zones. Devices, however, can belong to more than one zone. A *zone alias* is a collection of zone members. A zone alias can be added to one or more zones.

Note: In some cases, inactive zones might not have any zone members. Activating an inactive zone set that contains empty zones will fail if a switch does not support empty zones in active zone definitions.

Zone membership can be specified by:

- The N_Port_Name of the N_Port connected to the switch (also known as WWN zoning or port name zoning)
- The N_Port address identifier assigned during fabric login (also known as PortId or FCID zoning)
- The Node_Name associated with the N_Port
- The Domain identification (Domain_ID) and physical port ID of the Switch Port to which the N_Port is attached (also known as domain port zoning)
- An alias name

Zoning supports the use of aliases, which are meaningful names assigned to devices. An alias can also be a group of devices that are managed together to make zoning easier.

There are two types of zoning:

Hardware zoning (port zoning)

In hardware zoning (also called *port zoning*), the members of a zone are the physical ports on a fabric switch.

Software zoning (WWN zoning)

Software zoning uses the Simple Name Server (SNS) that runs inside a fabric switch. It is based on the node WWN or port WWN of the zone members to be included. Software zoning lets you create symbolic names for the zones and zone members.

A *default zone* is a group of devices that are not members of the active zone set. These can communicate with each other but not with members of any other zone. Default zoning is enabled by default. You can use a switch element manager to configure the Default Zone option to enable or disable the default zone independently of the active zone set.

Note:

1. If the default zone is disabled, devices that are not members of the active zone set cannot communicate.
2. If the default zone is disabled and no zone set is active, no devices can communicate.

3. If default zoning is enabled, deactivating the active zone set makes all devices members of the default zone. If default zoning is disabled, all communication stops.
4. If you activate one zone set while another zone set is active, the currently active zone set is deactivated.
5. If your EFC Manager manages multiple fabrics, ensure that you have the correct zone set for the fabric you are currently updating.

You can view zones in SANs using the topology viewer console. Fabric Manager lets you create, update, and delete zones, zone aliases, and zone sets in a SAN. In addition, you can do the following:

- Assign zones to zone sets
- Activate and deactivate zone sets
- View zone membership
- Add and remove zone members to and from a zone
- Add and remove zone members to and from a zone alias
- Assign zone aliases to zones

Note: Certain types of switches and the types of agents configured can affect the zoning operations that you can perform.

Zone configuration

When you configure zones, you can use Fabric Manager or the management application for the devices. The advantages of using Fabric Manager rather than the management applications of the devices are that Fabric Manager lets you perform zoning from a single interface, and you can use consistent, familiar methods to work with devices from multiple vendors.

Consider the following guidelines if you use Fabric Manager with zones:

- For Brocade and McDATA fabrics, topology information is collected through CIM agents. If no CIM agents are configured, or if the fabric is a QLogic or Cisco fabric, have at least one Storage Resource agent attached to each fabric and as a member of an active zone for the complete topology display. In this way, you can manage your switches and a string is returned, provided by the vendor, that can be the worldwide name (WWN) or a vendor, model, and level type designation.
- For zone discovery and zone control for any fabrics with Brocade switches or vendor Brocade switches, you must have a SMI-S agent configured for the fabric.
- A zone set can be deleted while it contains zones. With QLogic switches, those zones are put into an orphan zone set. The zones can later be moved from an orphan set into other zone sets. The Fabric Manager zone configuration panels cannot display zones in an orphan zone set. However, the zones exist on the switch, and they can be configured using the tool from the switch vendor.
- For McData switches, you can add empty zones to inactive zone sets. You cannot add empty zones to inactive zone sets for switches from other vendors.

Note: Activating an inactive zone set that contains empty zones can fail if a switch does not support empty zones in active zone definitions.

Note: Fabric probes through the Storage Resource agents do not detect orphan zones. These zones are not listed in the Zone Configuration window. However, the zones exist on the switch, and they can be configured using the tool from the switch vendor. However, Brocade orphan zones, which are discovered through

Storage Management Initiative (SMI) agents, are listed in the Zone Configuration window and can be configured by the Fabric Manager. Brocade orphan zone aliases are listed in the Zone Configuration window and can be configured by the Fabric Manager. Similarly, for McDATA fabrics managed where an SMI agent is configured, orphan zones are listed in the Zone Configuration window and can be configured by the Fabric Manager.

Switch zoning capabilities

When you select a fabric in which to perform zoning operations, Fabric Manager determines the capabilities of the switches in the SAN and limits the zoning operations based on that information.

To see a list of supported switches and their capabilities, default values, ranges, and possible effects, go to <http://www.ibm.com/support/docview.wss?uid=swg21386446>.

Zoning for Cisco MDS 9000 devices:

Cisco switches support virtual SANs (VSANs), which is the logical partitioning of a fabric into multiple fabrics. The overall network is referred to as the physical infrastructure, and the logical fabrics are the VSANs. Fabric Manager provides basic zone discovery and zone control for the Cisco MDS 9000 series of devices using the FC-GS-3 protocol. This allows in-band zone control and defaults to the Cisco VSANs.

The following tasks are supported:

- Creating, deleting, and updating zones
- Creating, deleting, updating, activating, and deactivating zone sets

Fabric Manager supports these zone member types:

- N_Port WWN
- FC ID (the fibre channel ID of an N_port attached to the switch)

Unsupported zone and zone member types are displayed as non-standard on the topology display. Fabric Manager does not allow zone management of zones and zone sets that contain unsupported members.

Brocade switches in zones:

Brocade switches support some non-standard zones such as quick loop zones, fabric assist zones, and protocol zones. If the switch configurations have these zones already defined, Fabric Manager preserves them and does not modify them in any way. They are not displayed on the zone control interface. However, if they are part of a zone set that is active at some time, the devices that are part of such zones that are also online are displayed in the topology Zone View.

You can create, change, and delete non-standard zones by using the Brocade switch management application.

Note:

- A zone set displayed on the zone control interface might have no zones. This could mean that the zone set contains a non-standard zone on a Brocade switch.
- A zone set might have non-standard zones on a Brocade switch that are not displayed on the zone control interface.

- When creating a new zone using the same zone name as that of a non-standard zone that already exists on a Brocade switch, you might get a error that the zone already exists. Because the non-standard zone is not displayed on the zone control interface, you must verify this by using the switch management application.
- Tivoli Storage Productivity Center allows you to add empty zones to an inactive zone definition. Activation of a zone set containing empty zones may still fail if the switch does not support empty zones in active zone definitions.

Agents

The IBM Tivoli Storage Productivity Center uses agents to gather data: Common Information Model (CIM) agents, Storage Resource agents, Data agents, Fabric agents, and out-of-band Fabric agents.

Common Information Model agents

Common Information Model (CIM) agents enable communication between the storage device (storage subsystem, fabric switch, or tape library) and IBM Tivoli Storage Productivity Center. Each storage device that you want to manage must have a CIM agent either installed or embedded unless it is a device that uses Native API connection.

CIM agents are provided by the vendor of the storage subsystem, fabric switch, or tape library. Each vendor provides unique agent code for their family of storage devices. This code implements a Common Information Model Object Manager (CIMOM) that conforms to the Storage Management Initiative Specification (SMI-S) of the Storage Networking Industry Association (SNIA).

The CIM agent enables communication between the storage device and Tivoli Storage Productivity Center. Commands and responses are transmitted between Tivoli Storage Productivity Center and the CIM agent using an XML transport layer. The CIM agent to storage device layer uses a vendor-specific proprietary interface.

The CIM agent usually must be installed and configured, so that it can identify the storage devices with which it communicates. Some storage devices, such as Cisco fabric switches, contain embedded CIM agents and so do not require that CIM agents be installed, although Tivoli Storage Productivity Center must be configured to point directly to the storage devices that contain the embedded CIM agents.

CIM agents can be referred to as SMI-S providers, CIM proxy agents, and CIMOM agents. CIM agents can be embedded in the device or installed on a separate computer.

Note:

- Do not install multiple CIM agents on a single computer because of port conflicts.
- Do not install a CIM agent on the system where a Tivoli Storage Productivity Center server component is installed.

Storage Resource agents and Data agents

Use Storage Resource agents and Data agents to collect information about the storage resources (for example, servers, computers, workstations, HBAs, and so on) in your environment. You must deploy an agent on systems that you want Tivoli Storage Productivity Center to monitor.

Note: In Tivoli Storage Productivity Center v4.2 and above, you can deploy Storage Resource agents only. If you want to install a Data agent, you must own a previous version of the product. See the documentation for the previous version of the product for information about how to install a Data agent.

Agents gather the following information about storage resources:

- Asset information
- File and file system attributes
- Database application information
- Novell Directory Services (NDS) tree information (only supported by Data agents)
- Network attached storage (NAS) device information
- Topology information
- Zoning information

Agents also discover fabrics and perform zone control for Cisco and Qlogic fabrics.

Deployment considerations for Storage Resource agents:

There are a number of guidelines and limitations that you must consider when managing the IBM Tivoli Storage Productivity Center agents in your environment.

Use the following information when deploying Storage Resource agents and migrating existing Data agents:

- Multiple Storage Resource agents that are probing or scanning the same storage resources
- Platforms that support Storage Resource agents deployments
- Product functions that are not available for storage devices monitored by Storage Resource agents
- Required authority for deploying Storage Resource agents
- Orphan zones
- Firewalls and Storage Resource agent deployments
- Multiple computers in a deployment
- Communication between the Tivoli Storage Productivity Center server and a Storage Resource agent
- Daemon and non-daemon services
- Authentication between the Tivoli Storage Productivity Center server and a Storage Resource agent
- Storage Resource agents and Data agents on the same computer
- Time zones for computers monitored by Storage Resource agents
- UNIX systems using RSH protocol
- Deployments on Windows - NetBIOS setting
- Deployments on Windows - User Account Control (UAC) remote restrictions
- Additional information about Storage Resource agents

Multiple Storage Resource agents that are probing or scanning the same storage resources

If multiple Storage Resource agents are set up to probe or scan the same storage resources, the Storage Resource agent that was added to Tivoli Storage Productivity Center first is used for the probe or scan. Therefore, only data that is gathered by the first Storage Resource agent is shown.

Platforms that support Storage Resource agents for the same

For a list of platforms on which you can deploy Storage Resource agents, go to the Tivoli Storage Productivity Center support site at http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center_Standard_Edition, click the **Documentation** link, and enter **Platform Support: Agents, Servers and GUI** in the **Search** box. Click the link to the document for the appropriate release.

Product functions that are not available for storage devices monitored by Storage Resource agents

Before you migrate an existing Data agent to a Storage Resource agent or deploy a new Storage Resource agent, make sure that the product functions you want to use on the monitored devices are available for those agents. The following functions are not available for storage devices that are monitored by Storage Resource agents:

- NetWare monitoring. You must use Data agents to monitor NetWare Servers.
- Monitoring of certain relational databases. For list of relational databases that can be monitored by Storage Resource agents, go to the Tivoli Storage Productivity Center support site at http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center_Standard_Edition, click the **Documentation** link, and enter **Platform Support: Agents, Servers and GUI** in the **Search** box. Click the link to the document for the appropriate release.
- Changing zone configuration and reporting of HBA, fabric topology, or zoning information for fabrics connected to hosts running Linux on IBM Power Systems™ or zSeries® hardware. These limitations also apply to Storage Resource agents on all guest operating systems for VMware configurations.

If you are already using these functions in Tivoli Storage Productivity Center, it is recommended that you do not migrate the corresponding Data agents to Storage Resource agent. If you plan to use these functions in the future, you must install Data agents on the computers that you want to monitor.

Required authorities for deploying and running Storage Resource agents

Before you can create deployment schedules and deploy Storage Resource agents on target computers, you must meet the following requirements:

- To create deployment schedules, you must be logged in to Tivoli Storage Productivity Center with a user ID that has the **Productivity Center Administrator** role. See the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* for more information about user roles.
- To deploy Storage Resource agents on target computers, you must provide a user ID that has administrative rights on those computers. You enter this ID when creating a deployment schedule. Tivoli Storage Productivity Center uses this ID to log in to the target computers and install and configure the necessary runtime files for the agents.

The user under which a Storage Resource agent (daemon or non-daemon) runs must have the following authorities on the target computers:

- On Linux or UNIX, the user must have root authority. By default, an agent runs under the user 'root'.

- On Windows, the user must have Administrator authority and be a member of the Administrators group. By default, a Storage Resource agent runs under the 'Local System' account.

Orphan zones

Storage Resource agents do not collect information about orphan zones. An orphan zone is a zone that does not belong to at least one zoneset.

Firewalls and Storage Resource agent deployments

Before you can deploy a Storage Resource agent on a computer, you must turn off the firewall on that computer. If you do not turn off the firewall, the deployment fails. To turn off the firewall on a Windows 2008 computer, complete the following steps:

1. On the desktop of the Windows 2008 computer, click **Start > Administrative Tools**.
2. Click **Windows Firewall with Advanced Security**.
3. Click **Windows Firewall Properties**.
4. Change the **Firewall state** field to **Off** on the following tabs: **Domain Profile, Private Profile, Public Profile**.
5. Click **OK** to accept the changes and exit.
6. Deploy a Storage Resource agent to the Windows 2008 computer.

Multiple computers in a deployment

You can define a deployment schedule to include multiple computers on which to install Storage Resource agents. The computers you include in a deployment schedule must share the same administrative user ID and password. IBM Tivoli Storage Productivity Center uses these user credentials to log in to the computers when installing Storage Resource agents. If the computers in a deployment schedule do not share the same administrative user credentials, you must create separate deployment jobs for them.

Note: When you deploy a Storage Resource agent to multiple computers, a globally unique identifier (GUID) is created for each computer (if one does not exist).

Communication between the Tivoli Storage Productivity Center server and a Storage Resource agent

The Tivoli Storage Productivity Center server connects to a monitored computer when a Storage Resource agent is deployed and whenever a data collection schedule runs against that agent.

During deployment, the server communicates with the target computer using one of the following protocols:

- Windows server message block protocol (SMB)
- Secure Shell protocol (SSH)
- Remote execution protocol (REXEC)
- Remote shell protocol (RSH)

After deployment, the type of communication between the server and agent on that computer depends on whether you deployed the agent as daemon service or non-daemon service. See the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* for more information about these protocols.

Daemon and non-daemon services

You can deploy a Storage Resource agent as a daemon or non-daemon service:

- A Storage Resource agent that is deployed as a daemon service runs in the background on the monitored computer and listens for requests from the Tivoli Storage Productivity Center server. Connectivity between the server and agent is established using SSL. The server and agent have their respective certificates and no additional information is required besides those certificates and the security that is provided by the SSL protocol.
- A Storage Resource agent deployed as a service on demand (non-daemon service) runs as a stand-alone executable file on the monitored computer. Communication from the server to the agent uses the same protocol that was used during the deployment of the agent. Communication from the agent to the server uses SSL.

Authentication between the Tivoli Storage Productivity Center server and a Storage Resource agent

Tivoli Storage Productivity Center requires the correct authentication information (user ID, password, port, certificate location, or passphrase) for monitored computers each time it communicates with Storage Resource agents on those computers. If the authentication information changes for a host computer on which a Storage Resource agent is deployed, the authentication information for that agent must be updated using the **Change Authentication** or **Update Storage Resource agent List** functions on the Data/Storage Resource agents panel located under the **Administrative Services > Data Sources** node in the navigation tree.

Storage Resource agents and Data agents on the same computer

You cannot deploy a Storage Resource agent on a computer where a Data agent is already installed *and* pointing to the same Data server as the Storage Resource agent. You can deploy a Storage Resource agent on the same computer as a Data agent if those agents communicate with different Data servers and use different ports when listening for requests.

Time zones for computers monitored by Storage Resource agents

The time zones of computers that are monitored by Storage Resource agents are shown as Greenwich Mean Time (GMT) offsets in Tivoli Storage Productivity Center reports. The time zones of computers that are monitored by Data agents are shown as the local time zone in Tivoli Storage Productivity Center reports. For example, a computer located in Los Angeles shows the following time zones in the By Computer report in Asset reporting.

- If monitored by a Storage Resource agent: (GMT-8:00) GMT-8:00
- If monitored by a Data agent: (GMT-8:00) America/Los_Angeles Time zone

UNIX systems using Remote Shell protocol (RSH)

If RSH is configured to use a user ID and password, the connection fails. To successfully connect to a system using RSH, you must set up the `.rhosts` file (in the home directory of the account). RSH must be configured to accept a login from the system that is running your application.

Deployments on Windows - NetBIOS setting

If you want to install a Storage Resource agent or Data agent on Windows targets, the **Enable NetBIOS over TCP/IP** option must be selected in the Control Panel settings for the computer's network connections properties.

To set this option, click **Start > Settings > Control Panel > Network and Dial-Up Connections > <some_connection> > Properties > Internet Protocol (TCP/IP) > Advanced > WINS > Enable NetBIOS over TCP/IP.**

See the documentation for your firewall to determine if these ports are not blocked for inbound requests.

To determine if security policies are blocking the connection ports, click **Start > Settings > Control Panel > Administrative Tools.**

Depending on whether your policies are stored locally or in Active Directory, follow these directions:

Locally stored policies

For locally stored policies, click **Start > Settings > Control Panel > Administrative Tools > Local Security Policy > IP Security Policies on Local Computer.**

Policies stored in Active Directory

For policies stored in Active Directory, examine the IP security policies and edit or remove filters that block the ports:

- Click **Administrative Tools > Default Domain Security Settings > IP Security Policies on Active Directory.**
- Click **Administrative Tools > Default Domain Controller Security Settings > IP Security Policies on Active Directory.**

For all Windows systems, the Server service must be running to connect to a Windows system using the Windows protocol.

The following table lists the ports reserved for NetBIOS. Ensure that these ports are not blocked.

Port	Description
135	NetBIOS Remote procedure call. (Not currently used.)
137	NetBIOS name service.
138	NetBIOS datagram. (Not currently used.)
139	NetBIOS session (for file and print sharing).
445	CIFS (on Windows XP).

For Windows Server 2008, shares must be shared for the Guest or Everyone accounts, and password protected sharing must be disabled. To disable password protected sharing, follow these steps:

1. Click **Control Panel > Networking and Sharing Center.**
2. Click the down arrow next to **Password protected sharing.**
3. Click **Turn off password protected sharing.**
4. Click **Apply.** Exit from the Control Panel.

Deployments on Windows 2008 - User Account Control (UAC) remote restrictions

If you are planning to install Storage Resource agents remotely on a Windows 2008 system, you must disable the User Account Control (UAC) remote restrictions on the Windows system. User Account Control is a security component on Windows.

Note: This task contains steps that tell you how to modify the registry. Serious problems might occur if you modify the registry incorrectly. Therefore, make sure that you follow these steps carefully. For added

protection, back up the registry before you modify it. Then you can restore the registry if problems occur. For information about how to back up and restore the registry, see <http://support.microsoft.com/kb/322756/>.

To disable UAC remote restrictions, follow these steps:

1. Click **Start > Run**. Enter **regedit** and click **OK**.
2. Locate and click the following registry subkey:
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\
Policies\System
3. If the LocalAccountTokenFilterPolicy registry entry does not exist, follow these steps:
 - a. On the **Edit** menu, click **New > DWORD Value**.
 - b. Enter **LocalAccountTokenFilterPolicy** as the name for the DWORD value in the content pane. Click **Enter**.
 - c. Right-click **LocalAccountTokenFilterPolicy**, then click **Modify**.
 - d. In the Edit DWORD Value window, enter **1**, then click **OK**.
This value can be 0 or 1:
 - 0 This value builds a filtered token. This value is the default value. The administrator credentials are removed.
 - 1 This value builds an elevated token.
 - e. Exit the registry editor.

Additional information about Storage Resource agents

See *Planning for Storage Resource agents* in the *IBM Tivoli Storage Productivity Center Installation and Configuration Guide* for more information about Storage Resource agents.

Inband fabric agents or Storage Resource agents

Inband fabric agents and Storage Resource agents collect information about the Storage Attached Network (SAN) fabrics. You must deploy agents on systems with host bus adapters (HBAs) that are connected to the SAN fabrics that you want to manage.

Note: In Tivoli Storage Productivity Center v4.2 and above, you can deploy Storage Resource agents only. If you want to install an Inband fabric agent (Fabric agent), you must own a previous version of the product. See the documentation for the previous version of the product for information about how to install a Fabric agent.

Agents communicate through the HBA and collect the following information:

- Attribute and topology information
- Host-level information, such as file system and volume mappings
- HBA information, including make, model, and driver versions
- Port state information
- Zoning information (for QLogic and Cisco fabrics, and for McDATA fabrics if SMI-S agents are not installed)

You can deploy Storage Resource agents locally or remotely from the Tivoli Storage Productivity Center server. Deploy a Storage Resource agent on *at least* one system that has an HBA that is connected to each switch that you want to manage; consider installing two Storage Resource agents for each switch to ensure redundancy.

Out-of-band fabric agents

Out-of-band fabric agents issue SNMP queries to the fabric switches and collect topology information.

Be sure that your out-of-band fabric agents point to each switch in the Storage Area Network (SAN) that you want to manage.

Out-of-band fabric agents are necessary for Virtual storage area network (VSAN) information for Cisco switches. For Brocade fabrics, if no SMI-S agent is configured, Out-of-band fabric agents are necessary for Zoning information for Brocade switches.

IBM Tivoli Storage Productivity Center only supports SNMP V1.

See the Planning > Planning for Fabric Manager > *Configuring Cisco switches* topic in the *Installation and Configuration Guide* for more information about using Out-of-band fabric agents and Cisco switches.

Tivoli Storage Productivity Center Monitoring Agent

The IBM Tivoli Storage Productivity Center Monitoring Agent can be used by the IBM Tivoli Enterprise Monitoring Server to monitor systems in your enterprise. This agent is an optional program you can install and use in your enterprise.

IBM Tivoli Monitoring monitors and manages system and network applications on a variety of operating systems, tracks the availability and performance of your enterprise system, and provides reports to track trends and troubleshoot problems. The Tivoli Storage Productivity Center Monitoring Agent participates in that environment.

For performance reasons, install the Tivoli Storage Productivity Center server on a separate system from the IBM Tivoli Monitoring server. You can install the Tivoli Storage Productivity Center Monitoring agent on any system you want to monitor.

Before using the Tivoli Storage Productivity Center Monitoring Agent, you must have IBM Tivoli Enterprise Monitoring Server installed. For information about the IBM Tivoli Enterprise Monitoring Server, see http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc_6.2.1/welcome.htm.

For information about installing IBM Tivoli Monitoring, see the Quick Start Guide at <http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc/GI11-8058-00.htm>.

To use the Tivoli Storage Productivity Center Monitoring Agent, install this agent on the system you want to monitor.

To see the type of reports displayed by the Tivoli Storage Productivity Center Monitoring Agent, see “Tivoli Storage Productivity Center Monitoring Agent reports” on page 1130.

Security

The IBM Tivoli Storage Productivity Center environment supports role-based security. Tivoli Storage Productivity Center also has several applications that use certificates to ensure server to server security.

For role-based authorization, user groups contained in the authentication repository, which is either the local operating system or an LDAP-compliant directory depending on the choice of authentication mechanism made during the Tivoli Storage Productivity Center installation, are associated with predefined roles. When a user ID is authenticated to Tivoli Storage Productivity Center through the GUI, CLI, or APIs, membership in a user group determines the authorization level.

If you are using a previous version of Tivoli Storage Productivity Center (before v4.2) or are still using Data agents in v4.2, Agent Manager includes demonstration certificates that you use for a test environment, but you can also create new generated certificates during the installation of the Agent Manager. These generated certificates are used when the Tivoli Storage Productivity Center server or agent registers with the Agent Manager.

Role-based authorization

Operating system groups or LDAP groups (for example, groups contained in your LDAP-compliant repository) are associated with predefined roles. When a user ID is authenticated to Tivoli Storage Productivity Center through the GUI, CLI, or APIs, membership in an operating system or LDAP group determines the authorization level of the user.

Table 9 shows the association between roles and authorization level.

If you select OS authentication for your Tivoli Storage Productivity Center installation, then you do not have to create any of the groups before installation. The Tivoli Storage Productivity Center superuser role automatically gets mapped to the Administrators group on Windows, to the system group on AIX, or to the root group on Linux.

Note: If you plan to select LDAP authentication during your Tivoli Storage Productivity Center installation, then the group you intend to map to the Tivoli Storage Productivity Center Superuser role must exist in the LDAP-compliant directory before you start your Tivoli Storage Productivity Center installation.

Table 9. Roles and authorization levels

Role	Authorization level
Superuser	Has full access to all Tivoli Storage Productivity Center functions.
Productivity Center administrator	Has full access to operations in the Administration section of the GUI.
Disk administrator	Has full access to Tivoli Storage Productivity Center disk functions.
Disk operator	Has access to reports only for Tivoli Storage Productivity Center disk functions. These reports include information about tape devices.
Fabric administrator	Has full access to Tivoli Storage Productivity Center Fabric Manager functions.
Fabric operator	Has access to reports only for Tivoli Storage Productivity Center Fabric Manager functions.
Data administrator	Has full access to Tivoli Storage Productivity Center Data Manager functions.
Data operator	Has access to reports only Tivoli Storage Productivity Center Data Manager functions.

Table 9. Roles and authorization levels (continued)

Role	Authorization level
Tape administrator	Has full access to Tivoli Storage Productivity Center tape functions.
Tape operator	Has access to reports only for tape functions.

Notes:

1. If a user has multiple roles, the authorization level is a combination of the levels for each of the roles.
2. If a user is not a member of any of the roles listed, no access is granted to that user.
3. For enterprise-rollup reports, you need superuser or Tivoli Storage Productivity Center administrator authority to do the following actions:
 - Add, remove, or modify the Tivoli Storage Productivity Center subordinate server that the master server is monitoring.
 - Create or run probe jobs that include Tivoli Storage Productivity Center subordinate servers.

Any Tivoli Storage Productivity Center user can generate enterprise-rollup reports.

4. When you create and save role-to-group mappings in Tivoli Storage Productivity Center, these mappings get propagated into Tivoli Integrated Portal, where the groups are given the operator authorization. Occasionally, you might find that after creating and saving the role-to-group mappings in Tivoli Storage Productivity Center, you are unable to access Tivoli Integrated Portal as a valid user (in a valid group that is mapped to the operator authorization in Tivoli Integrated Portal). If this condition occurs, stop and restart the Tivoli Integrated Portal server.

Agent Manager certificates

Agent Manager provides a demonstration certificate to facilitate the rapid deployment of a test environment or demonstration environment.

The demonstration certificate and its password are publicly available. Anyone who has a demonstration certificate can participate in your deployment. Using the demonstration certificate does not provide the level of security required by a typical production environment. In a production environment, let the Agent Manager installation program generate a certificate that is unique to Agent Manager installation. When you install Agent Manager, you can specify to use the demonstration certificates or create a new certificate on the *Specify Certificates* page.

Note: You cannot install Agent Manager in Tivoli Storage Productivity Center v4.2 and above. If you want to install Agent Manager, you must use a previous version of the product.

SAN Volume Controller topics

This section contains information about the following SAN Volume Controller topics: managed disks (MDisks), managed-disk states, managed-disk groups, and virtual disks (VDisks).

IBM Tivoli Storage Productivity Center SAN Volume Controller

The IBM Tivoli Storage Productivity Center SAN Volume Controller is a SAN appliance that attaches open-systems storage devices to supported open-systems hosts.

The SAN Volume Controller provides symmetric virtualization by creating a pool of managed disks from the attached storage subsystems, which are then mapped to a set of virtual disks for use by attached host computer systems. System administrators can view and access a common pool of storage on the SAN, which enables them to use storage resources more efficiently and provides a common base for advanced functions. The SAN Volume Controller performs the following functions for the SAN storage that it controls:

- Creates a single pool of storage
- Manages logical volumes
- Provides advanced functions for the SAN, such as large scalable cache, copy services, and space management

A *node* is a single storage engine. The storage engines are always installed in pairs with one or as many as four pairs of nodes constituting a *cluster*. Each node in a pair is configured to back up the other.

Each pair of nodes is known as an *I/O group*. All I/O operations are cached on both nodes for resilience. To eliminate any single point of failure, each of the two nodes are protected by different uninterruptible power supplies. A node can be in only one I/O group.




Figure 4. A SAN Volume Controller node

The SAN Volume Controller I/O groups see the storage presented to the SAN by the back-end controllers as a number of disks, known as *managed disks*. The application services do not see these managed disks. Instead, they see a number of logical disks, known as *virtual disks*, that are presented to the SAN by the SAN Volume Controller. Each node must only be in one I/O group and provides access to the virtual disks in the I/O group.

The fabric contains two distinct zones: a host zone and a disk zone. In the host zone, the host systems can identify and address the nodes. You can have more than one host zone. Generally, you create one host zone per operating system type. In the disk zone, the nodes can identify the disk drives (storage subsystems). Host systems cannot operate on the disk drives directly; all data transfer occurs through the nodes. As shown in Figure 5, several host systems can be connected to a SAN fabric. A cluster of SAN Volume Controllers is connected to the same fabric and presents virtual disks to the host systems. You configure these virtual disks using the disks located on the RAID controllers.

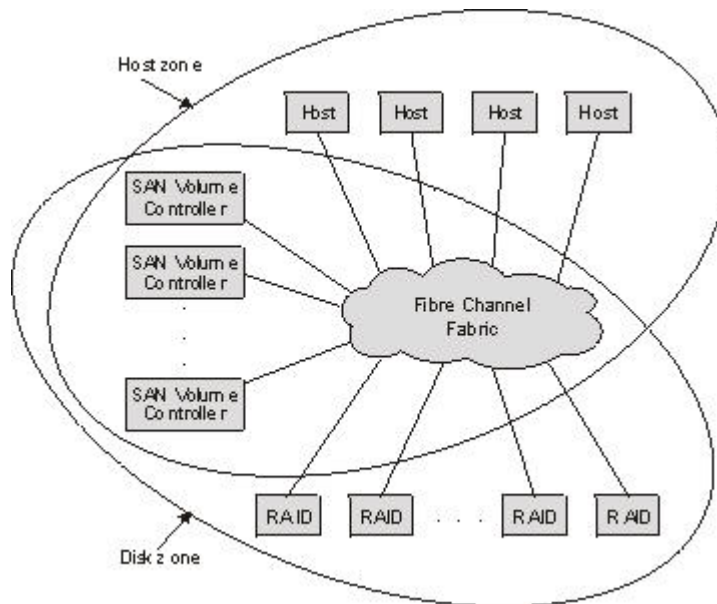


Figure 5. Example of a SAN Volume Controller in a fabric

Managed disks

A managed disk (MDisk) is a logical disk that a SAN Volume Controller has exported to the SAN fabric to which the nodes in the cluster are attached.

A managed disk might consist of multiple physical disks that are presented to the SAN as a single logical disk. A managed disk always provides usable blocks of physical storage to the cluster, even if it does not have a one-to-one correspondence with a physical disk.

Figure 6 on page 131 shows physical disks and managed disks.




Figure 6. Physical disks and managed disks (MDisks)

Extents

Each MDisk is divided into chunks of equal size called *extents*. Extents are numbered sequentially (beginning at 0) from the start to the end of the managed disk. Extents are a unit of mapping the data between MDisks and virtual disks (VDisks).

The extent size is a property of managed-disk groups (MDisk groups). When a managed disk is added to a managed-disk group, the size of the extents that the MDisk will be broken into depends on the attribute of the managed-disk group to which it has been added.

Access modes

The access mode determines how the SAN Volume Controller cluster uses the MDisk. The possible modes are the following:

Unmanaged

The MDisk is not used by the cluster.

Managed

The MDisk is assigned to an MDisk group and provides extents that virtual disks (VDisks) can use.

Image The MDisk is assigned directly to a VDisk with a one-to-one mapping of extents between the MDisk and the VDisk.

Managed disk states

This topic describes the operational states of managed disks (MDisks).

Managed disks can be in the following states:

Online

The MDisk can be accessed by all online nodes. That is, all the nodes that are currently working members of the cluster can access this MDisk. The MDisk is online when the following conditions are met:

- All timeout error recovery procedures complete and report the disk as online.
- Volume inventory of the target ports correctly reported the MDisk.
- Discovery of this volume completed successfully.

- All of the managed disk target ports report this volume as available with no fault conditions.

Degraded

The MDisk cannot be accessed by all the online nodes. That is, one or more (but not all) of the nodes that are currently working members of the cluster cannot access this MDisk. The MDisk might be partially excluded; that is, some of the paths to the MDisk (but not all) have been excluded.

Excluded

The MDisk has been excluded from use by the cluster after repeated access errors.

Offline

The MDisk cannot be accessed by any of the online nodes. That is, all of the nodes that are currently working members of the cluster cannot access this MDisk. This state can be caused by a failure in the SAN, the storage subsystem, or one or more physical disks connected to the storage subsystem. The MDisk will only be reported as offline if all paths to the disk fail.

Attention: If your fabric is undergoing transient link breaks or you have been replacing cables or connections in your fabric, you might see one or more MDisks change to the degraded status. If an I/O operation was attempted during the link breaks and the same I/O failed several times, the MDisk will be partially excluded and change to a status of degraded. You must include the MDisk to resolve the problem.

Managed disk path Each managed disk will have an online path count, which is the number of nodes that have access to that managed disk; this represents a summary of the I/O path status between the cluster nodes and the particular storage device. The maximum path count is the maximum number of paths that have been detected by the cluster at any point in the past. Thus if the current path count is not equal to the maximum path count then the particular managed disk may be degraded. That is, one or more nodes may not see the managed disk on the fabric.

Managed disk groups

A SAN Volume Controller *MDisk group* is a collection of MDisks that jointly contain all the data for a specified set of virtual disks (VDisks).

All MDisks in a group are split into extents of the same size. VDIs are created from the extents that are available in the group. You can add MDisks to an MDisk group at any time in order to increase the number of extents that are available for new VDIs or to expand existing VDIs. You can add only MDisks that are in unmanaged mode. When MDisks are added to a group, their mode changes from unmanaged to managed.

Figure 7 on page 133 shows an MDisk group containing four MDisks.




Figure 7. MDisk group

You can delete MDisks from a group under the following conditions:

- VDIs are not using any of the extents that are on the MDisk.
- Enough free extents are available elsewhere in the group to move any extents that are in use from this MDisk.

Attention: If you delete an MDisk group, you destroy all the VDIs that are made from the extents that are in the group. If the group is deleted, you cannot recover the mapping that existed between extents that are in the group and the extents that VDIs use. The MDisks that were in the group are returned to unmanaged mode and can be added to other groups. Because the deletion of a group can cause a loss of data, you must force the deletion if VDIs are associated with it.

Table 10 describes the operational states of an MDisk group.

Table 10. Managed disk group states

Status	Description
Online	The MDisk group is online and available. All the MDisks in the group are available.
Degraded	The MDisk group is available; however, one or more nodes cannot access all the MDisks in the group.
Offline	The MDisk group is offline and unavailable. No nodes in the cluster can access the MDisks. The most likely cause is that one or more MDisks are offline or excluded.

Attention: If a single MDisk in an MDisk group is offline and therefore cannot be seen by any of the online nodes in the cluster, then the MDisk group of which this MDisk is a member goes offline. This causes *all* the VDIs that are being presented by this MDisk group to go offline. Care should be taken when creating MDisk groups to ensure an optimal configuration.

Consider the following guidelines when you create MDisk groups:

- If you are creating image-mode VDIs, do not put all of these VDIs into one MDisk group because a single MDisk failure results in all of these VDIs going offline. Allocate your image-mode VDIs between your MDisk groups.
- Ensure that all MDisks that are allocated to a single MDisk group are the same RAID type. This ensures that a single failure of a physical disk in the storage subsystem does not take the entire group offline. For example, if you have three RAID-5 arrays in one group and add a non-RAID disk to this group, then you lose access to all the data striped across the group if the non-RAID disk fails. Similarly, for performance reasons you should not mix RAID types. The performance of all VDIs will be reduced to the lowest performer in the group.
- If you intend to keep the virtual disk allocation within the storage exported by a storage subsystem, ensure that the MDisk group that corresponds with a single subsystem is presented by that subsystem. This also enables nondisruptive migration of data from one subsystem to another subsystem and simplifies the decommissioning process if you want to decommission a controller at a later time.

Extents

To track the space that is available, the SAN Volume Controller divides each MDisk in an MDisk group into chunks of equal size. These chunks are called *extents*, and are indexed internally. Extent sizes can be 16, 32, 64, 128, 256, or 512 MB.

You must specify the extent size when you create a new MDisk group. You cannot change the extent size later; it must remain constant throughout the lifetime of the MDisk group. MDisk groups can have different extent sizes. However, different extent sizes can place restrictions on the use of data migration. The choice of extent size affects the total amount of storage that can be managed by a cluster. Table 11 shows the maximum amount of storage that can be managed by a cluster for each extent size. Because the SAN Volume Controller allocates a whole number of extents to each virtual disk that is created, using a larger extent size can increase the amount of wasted storage at the end of each virtual disk. Larger extent sizes also reduce the ability of the SAN Volume Controller to distribute sequential I/O workloads across many managed disks. Therefore, larger extent sizes might reduce the performance benefits of virtualization.

Table 11. Capacities of the cluster given extent size

Extent size	Maximum storage capacity of cluster
16 MB	64 TB
32 MB	128 TB
64 MB	256 TB
128 MB	512 TB
256 MB	1 PB
512 MB	2 PB

Virtual disks

An SAN Volume Controller *VDisk* is a logical disk that the cluster presents to the storage area network (SAN).

Application servers on the SAN access *VDisks*, not managed disks (*MDisks*). *VDisks* are created from a set of extents in an *MDisk* group. There are three types of *VDisks*: striped, sequential, and image.

Types

You can create the following types of *VDisks*:

Striped

The striping is at extent level. One extent is allocated, in turn, from each managed disk that is in the group. For example, a managed disk group that has 10 *MDisks* takes one extent from each managed disk. The 11th extent is taken from the first managed disk, and so on. This procedure, known as a round-robin, is similar to RAID-0 striping.

You can also supply a list of *MDisks* to use as the stripe set. This list can contain two or more *MDisks* from the managed disk group. The round-robin procedure is used across the specified stripe set.

Attention: Care should be taken when specifying a stripe set if your *MDisk* group contains *MDisks* of unequal size. By default striped *VDisks* are striped across all *MDisks* in the group. If some of the *MDisks* are smaller than others, the extents on the smaller *MDisks* will be used up before the larger *MDisks* run out of extents. Manually specifying the stripe set in this case might result in the *VDisk* not being created.

Figure 8 shows an example of a managed disk group containing three *MDisks*. This figure also shows a striped virtual disk created from the extents available in the group.




Figure 8. Managed disk groups and *VDisks*

Sequential

When selected, extents are allocated sequentially on one managed disk to create the virtual disk if enough consecutive free extents are available on the chosen managed disk.

Image Image-mode *VDisks* are special *VDisks* that have a direct relationship with one managed disk. If you have a managed disk that contains data that you want to merge into the cluster, you can create an image-mode virtual disk. When you create an image-mode virtual disk, a direct mapping is made between extents that are on the managed disk and extents that are on the virtual disk. The managed disk is not virtualized. In other words, the logical block address (LBA) *x* on the managed disk is the same as LBA *x* on the virtual disk.

When you create an image-mode VDisk, you must assign it to a managed disk group. An image-mode VDisk must be at least one extent in size. In other words, the minimum size of an image-mode VDisk is the extent size of the MDisk group to which it is assigned.

The extents are managed in the same way as other VDIsks. When the extents have been created, you can move the data onto other MDisks that are in the group without losing access to the data. After you move one or more extents, the virtual disk becomes a real virtualized disk, and the mode of the managed disk changes from image to managed.

Attention: If you add an MDisk to an MDisk group as a managed disk, any data on the MDisk will be lost. Ensure that you create image-mode VDIsks from the MDisks that contain data before you start adding any MDisks to groups.

MDisks that contain existing data have an initial mode of unmanaged, and the cluster cannot determine whether it contains partitions or data.

A virtual disk can have one of three states. Table 12 describes the different states of a virtual disk:

Table 12. Virtual disk states

Status	Description
Online	The virtual disk is online and available if both nodes in the I/O group can access the virtual disk. A single node will only be able to access a VDisk if it can access all the MDisks in the MDisk group associated with the VDisk.
Offline	The VDisk is offline and unavailable if both nodes in the I/O group are missing or none of the nodes in the I/O group that are present can access the VDisk.
Degraded	The status of the virtual disk is degraded if one node in the I/O group is online and the other node is either missing or cannot access the virtual disk.

If you delete a virtual disk, you destroy access to the data that is on the virtual disk. The extents that were used in the virtual disk are returned to the pool of free extents that is in the managed disk group. The deletion might fail if the virtual disk is still mapped to hosts. The deletion might also fail if the virtual disk is still part of a FlashCopy or a Metro Mirror mapping. If the deletion fails, you can specify the force-delete flag to delete both the virtual disk and the associated mappings to hosts. Forcing the deletion will also delete the copy services relationship and mappings.

Scenarios

This section contains scenarios which illustrate how to perform specific tasks using IBM Tivoli Storage Productivity Center.

The provided scenarios set out situations that can be solved using Tivoli Storage Productivity Center. When using these scenarios you should remember the following:

- The scenarios are intended to be step-by-step instructions that you follow to complete a specific job. Each scenario contains multiple tasks that must be followed in the order given.
- Specific values are given in the scenarios, such as user name, IP Address, probe name, and so forth. These are for illustration purposes only and you must replace them with values appropriate for your system.
- The scenarios do not provide in-depth conceptual information about the tasks. See the appropriate topics in the IBM Tivoli Storage Productivity Center Information Center for more information about the concepts behind the tasks.
- You will complete a scenario successfully if you follow the instructions correctly and your system performs as expected.

Monitoring and reporting on I/O utilization

This scenario involves monitoring a DS8000 and a DS4000 to determine what data can be moved to the less expensive storage media. Monitoring over time collects the necessary information for analysis. Trending reports, produced with IBM Tivoli Storage Productivity Center Reporter, display utilization information and can be used by management to make decisions about future storage needs.

You have a DS8000 and a DS4000. The DS8000 is nearly full but the DS4000 is only half full. You need to move workload onto the DS4000 as the demand for the DS8000 grows. You will monitor the utilization of volumes on each system to identify potential migration candidates. You will also gather this information over time to indicate trends. This trending information, in report format, will be used for future planning and will be presented to management who will make decisions based on the report.

To perform this scenario, complete the following steps:

Perform performance monitoring

1. **Configure the storage systems for monitoring by Tivoli Storage Productivity Center** This process is in two parts as Tivoli Storage Productivity Center uses native interfaces to connect to DS8000 storage systems while the DS4000 must be connected using a CIM object manager (CIMOM) connection. Use the Configure Devices wizard to configure the DS8000 system. For the DS4000 system, you can configure it manually by running a first discovery to locate the CIMOMs that are visible to the Device server with SLP. After configuring the CIMOM login information, run a second discovery to gather information about the DS4000 system. Run a CIMOM discovery every four hours. See “Discover storage subsystem, tape library, and fabric information with CIM agents” on page 330 for more information. Alternatively, you can configure the DS4000 system by using the Configure Devices wizard.

Configure the DS8000 system

- a. Use the Configure Devices wizard to enter connection information for the native interface of the DS8000 system. See “Configure storage subsystems” on page 306 for details of how to use the Configure Devices wizard. You can use the wizard to include the DS8000 system automatically in a probe. If you want to monitor SAN Volume Controller systems or the XIV systems, you can also use the Configure Devices wizard to enter connection information for their native interfaces.

Configure the DS4000 system manually

- a. Expand **Administrative Services > Data Sources > CIMOM Agents** .

- b. For each CIMOM you use, click the magnifying glass to the left of the CIMOM name to open the View and Edit window. Enter the name and password information. Check the **Test CIMOM connectivity before editing**.
- c. Click the **Save** icon. This action saves your information and performs a connectivity check. You must use this process for each SLP-discovered CIMOM you choose to use.
- d. Expand **Administrative Services > Discovery**, right-click **CIMOM**, and click **Run Now**.
- e. In the **CIMOM Discovery job submitted** window, click **OK**.
- f. When the job has finished, you can view the list of discovered storage systems by expanding **Disk Manager > Storage Subsystems**.

Configure the DS4000 system using the Configure Devices wizard

- a. Use the Configure Devices wizard to enter connection information for the native interface of the DS8000 system. See “Configure storage subsystems” on page 306 for details of how to use the Configure Devices wizard.
2. **Create a probe** Create the probe to run during off-hours.
 - a. Expand **Tivoli Storage Productivity Center > Monitoring > Probes**. Right-click **Probes** and select **Create Probe**.
 - b. In the **What to Probe** tab, move the DS6000 and the DS4000 from the **Available** list box to the **Current Selections** box.
 - c. In the **When to run** tab you will schedule when you want the probe to run. We strongly recommend that you run probes when nothing else is running.
 - 1) Choose to run the job repeatedly according to a schedule you specify.
 - 2) Select the current day, month, and year as the date when you want data collection to begin. Select 1:00 AM as the time the job should begin running.
 - 3) Select the days the job should run. Select all of them. A check mark will appear next to each day you have chosen.
 - d. In the **Alerts** tab accept the default (error log) as the means of notification if the job fails.
 - e. Click **Save**. The probe will run daily at 1:00 AM.
 3. **Create a storage-system performance monitor** Create a storage-system performance monitor to gather performance data and view the results in the log file.
 - a. Expand **Disk Manager > Monitoring**, right-click **Subsystem Performance Monitor** and select **Create Performance Monitor**.
 - b. In the Create Storage Subsystem Performance Monitor window, click the Storage Subsystems tab. Select the DS8000 and DS4000 you want to report on and move them to the Current Selections field. Make sure these storage systems have been probed. If they have not been, data collection will fail.
 - c. Click the Sampling and Scheduling tab. In the Sampling area, specify the length of the sampling window and how frequently the performance data is to be gathered, using the following values:
 - Interval length = 5 minutes
 - Duration = Indefinite
 - Scheduling = Begin immediately
 - d. In the Alert tab, **Monitor Failed** is the only triggering condition. In the Triggered Actions area accept the default (error log) means of notification.
 - e. Click **Save**. When prompted, type a name for the threshold and click **OK**. The name of the storage system performance monitor is displayed under

the Subsystem Performance Monitors node. The performance monitor will take 3 hours to run, as set in step c. The status icon will indicate when the job has completed successfully.

- f. When the job has completed, a table is displayed listing the storage systems. To view the log file for each storage system, select the magnifying glass to the left of the storage systems.

4. Tivoli Storage Productivity Center Reporter for Disk

This scenario uses Tivoli Storage Productivity Center Reporter for Disk to prepare a finished report automatically. Tivoli Storage Productivity Center Reporter is a free Java 2 Platform application that extracts storage system information and hourly performance statistics. The statistics are transcribed into a paper-style PDF file containing your storage server utilization, which is saved on the local machine.

If you do not have Tivoli Storage Productivity Center Reporter for Disk, complete the following steps. If you have installed Tivoli Storage Productivity Center Reporter for Disk, go to step 5.

- a. Go to the Tivoli Storage Productivity Center Reporter for Disk web page, <http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/PRS2618>.
- b. Select **Download Tivoli Storage Productivity Center Reporter for Disk for Windows (60MB)**
- c. Follow steps to download Tivoli Storage Productivity Center Reporter.
- d. Run the Installer

5. Create your report with Tivoli Storage Productivity Center Reporter

- a. To start the program go to **Start > Programs > Tivoli Storage Productivity Center Reporter for Disk > Tivoli Storage Productivity Center Reporter for Disk**. Select **Start**. The Confirmation window will open.
- b. Enter the Hostname (or IP address), Port, UserID, and Password of the system that you want the report on. Click **Continue**. The Select Serial Number window will open. Tivoli Storage Productivity Center Reporter for Disk saves a list of the commands and the status of the commands you are using in the GUI. The commands and status of the commands are displayed in the Status window.
- c. If the probe of the selected systems was successful, the DS8000 and DS4000 storage systems will be listed in the Select Serial Number window. Select the DS8000 storage system. Click **Select**. You can only create a report for one storage system at a time.
- d. The Select Date Range window opens. Select the **Last 7 days** radio button. Click **Select**,
- e. The Customer Information window opens. This information will be included in the title information of the finished report. The first text field, **Company**, is required. This is the company name that will appear on the title page of the report. All other fields are optional. When finished, click **Continue**.
- f. The IBM/Partner Information window opens. All fields are optional. Click **Continue**.
- g. Select a file name and directory in which to save the generated report. The default directory is "My Documents" and the default file name is TPCReport.pdf. It is strongly recommended that you save the PDF file under a different name. Click **Save**.

- h. The steps Tivoli Storage Productivity Center Reporter takes to create the report are shown in the status window. Once the report has been generated, the **Exit** button becomes available. Click **Exit**.
 - i. Go to "My Documents", find and double-click the PDF report.
The report includes the following information for storage systems, ports, arrays, and volumes:
 - General information
 - Performance summary
 - Statistics
 - Definitions
6. Using the DS8000 report, look at the Volume Performance Summary for volumes with low levels of I/O and throughput in Megabytes. These might be candidates for moving to the DS4000.

Printing

The report will have been saved as a PDF file using the latest Acrobat Reader on your system. If you only have the Reader of that version of Acrobat, you will be able to view the report, but you will be unable to print it. In this case take the following steps to print the report:

1. Open your complete version of Adobe Acrobat.
2. Select **File > Open** and navigate to the report.
3. Select **File > Print** to print the report.
4. Close the report. You will not be able to save the report as an earlier version of an Acrobat PDF file.

The following documents were referenced for this scenario:

- *IBM Tivoli Storage Productivity Center User's Guide GC32-1775*
- *Tivoli Storage Productivity Center Reporter for Disk Version 1.01 - Process and Report Overview*

Storage provisioning using workflows

IBM Tivoli Storage Productivity Center and Tivoli Provisioning Manager Workflow automation can reduce provisioning time from hours to minutes and significantly simplify the task of provisioning new storage capacity. This automation is especially valuable when adding storage provisioning to existing automated server provisioning tasks, such as deploying operating systems, software and network resources. This scenario shows you how to use Tivoli Provisioning Manager and Tivoli Storage Productivity Center Workflows to automate provisioning of SAN attached storage capacity to a server in a Microsoft Windows environment.

Tivoli Storage Productivity Center and Tivoli Provisioning Manager work together to provide automated storage provisioning for the data center.

Tivoli Provisioning Manager provides the means to automate common tasks in the data center that surround each of the four major pillars of provisioning:

- OS provisioning (Windows, Linux, AIX, Solaris, HP-UX, and so forth)
- Software provisioning (DB2, Web Sphere, Oracle, and many more)
- Network provisioning (provisioning VLANs, ACLs, IP addresses, activating and deactivating switch ports, and so forth)
- Storage provisioning (attaching SAN disks for use by the server)

Tivoli Provisioning Manager Workflows automate tasks, especially those that are complex, repetitive, and prone to error by applying best practice rules for zoning, device configuration, and path selection automatically.

Tivoli Storage Productivity Center provides heterogeneous management functions for the SAN storage environment. Tivoli Storage Productivity Center Workflows provide current status information on its SAN storage environment to Tivoli Provisioning Manager. They also allow Tivoli Provisioning Manager to make changes to the SAN through Tivoli Storage Productivity Center.

The following illustration shows the movement of workflow information between Tivoli Provisioning Manager and Tivoli Storage Productivity Center.

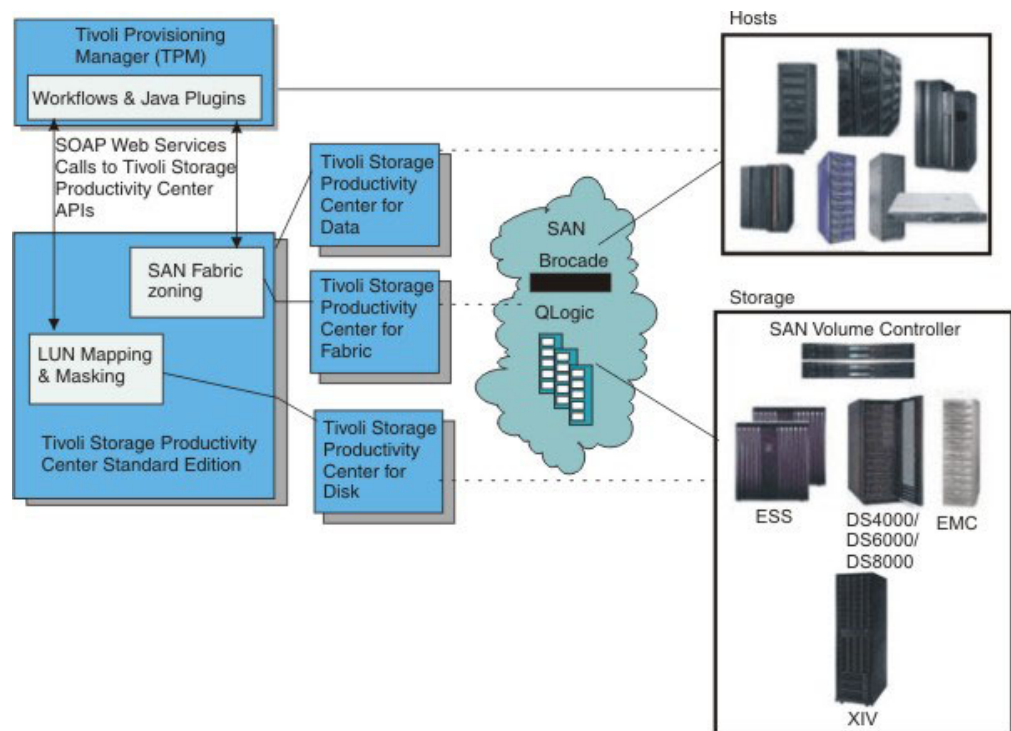


Figure 9. High-level architecture

End-to-end provisioning of SAN storage is difficult due to the complexity of modern environments and inputs that need to be considered. This scenario creates a Tivoli Provisioning Manager storage template used for end-to-end storage provisioning that can automate the process. Once the storage template has been created, it can be re-used to automate end-to-end provisioning of SAN storage with a few mouse clicks.

Without workflows you would need to perform the following individual steps in Tivoli Storage Productivity Center to provision storage:

- Add a volume (storage subsystem)
 - Select storage subsystem
 - Select or create new volume
 - Select host HBA ports (WWNs)
 - Select subsystem controller ports (WWNs)
 - Map volume to controller ports
 - Map volume to host HBA ports
- Set paths (SAN fabric switches)

- Determine if multiple paths are required
- Create or update zone(s)
- Get active zone set
- Add zone to zone set
- Activate zone set
- Map the HBA Volumes to the operating system and file system (Host server)
 - Rediscover attached volumes
 - Configure multipathing drivers
 - Map volumes to OS devices
- Create volume and file system (Host server)
 - Create Windows volume
 - Create or extend NTFS file system

With workflows, once the template has been created, these tasks are automated and condensed to the following:

- Select the host that needs additional storage
 - Select the template name from the template list
 - Click OK.

In order to use workflows, the following prerequisites must be met:

- Tivoli Provisioning Manager v5.1 has been installed and configured.
- Tivoli Storage Productivity Center v4.1 has been installed and configured.
- The following two automation packages have been downloaded, installed, and configured from the Open Process Automation Library (OPAL):
 - IBMTPC.tcdriver
 - WindowsNTFS.tcdriver

For more information on the prerequisites needed to use workflows, see the IBM White Paper, *Tivoli Provisioning Manager and IBM Tivoli Storage Productivity Center-Storage Provisioning Configuration and Exercises*.

To create and use a storage template, complete the following steps:

1. To create the Tivoli Provisioning Manager workflow template, complete the following steps:
 - a. Select **Inventory > Manage Templates > Storage Templates** from the Tivoli Provisioning Manager navigation tree.
 - b. Select the **Edit** button and click **Add Storage Template. Create Storage Template** window opens.
 - c. Enter **NTFS_SVC** as the name for the storage template. Click **Next**.
 - d. Define the volume container settings. In the **Configure Volume Container Settings** window, enter the following information:
 - Type **NTFS** for the **Volume Container Name**.
 - Select **NTFS** as the **Storage Manager**.
 - Click the **Add Volume Container Settings** button;
 - Click **Next**.
 - e. Define the physical volume settings
 - Scroll down to the **SAN Physical Volume Settings** and specify a SAN volume.
 - Type **SAN** in the **Name** field.
 - In the **Volume Container Settings** field, select **NTFS**.

- Select **mdiskgrp0** as the appropriate storage pool from **Storage Pool/Storage Subsystem**.
 - Type **5G** as the value for **Consumable Size**.
 - Set the **Function Type** to **Generic**.
 - For the **RAID Level**, select **UNKNOWN**.
 - Click **Add SAN Setting** and click **Next**.
- f. Define the data path settings:
- In the **Configure Data Path Settings** window, select **SAN** from the **SAN Physical Volume Setting** list box.
 - Click **Add Data Path Settings** and click **Next**.
- g. Define the logical volume settings:
- In the **Configure Logical Volume Settings** window, type **TPMVolume** in the **Name** field.
 - Select **Simple** as the **Logical Volume Type**.
 - Select **NTFS** for the **Volume Container Settings**.
 - Specify **5G** for **Consumable Size**.
 - Set the **Function Type** to **Generic**.
 - Set the **RAID Level** to **UNKNOWN**.
 - Click **Add Logical Volume Settings** and click **Next**.
- h. Define file system settings:
- In the **Configure File System Settings** window, select **TPMVolume** from the **Logical Volume Settings** list.
 - Type **H:** as the **Label**.
 - For the **File System Type**, select **NTFS** from the list.
 - Specify **5G** for **File System Size**.
 - Click **Add File System Settings** and click **Next**.
- i. File system mount settings are not used. In the **Configure Filesystem Mount Settings** window, click **Next**.
- j. Define disk partition settings:
- In the **Configure Disk Partition Settings** window, set the **Partition Size** to **5G**.
 - From the **Physical Volume Settings** list, select **SAN**
 - From the **Logical Volume Settings** list, select **TPMVolume**.
 - Click **Add Disk Partition Settings** and click **Next**.
- k. Verify the summary information. When you are done, click **Finish**.
- l. Verify that the template, **NTFS_SVC**, appears in the Tivoli Provisioning Manager GUI at **Inventory > Manage Templates > Storage Templates**.

This template will allow any system administrator to perform end to end provisioning.

2. To perform provisioning, complete the following steps
 - a. From the **General** tab of a Windows computer in Tivoli Provisioning Manager, click the **Edit** button and select **Add Storage to Host**.
 - b. Select **NTFS_SVC** from the **Storage Template** list
 - c. Click **OK**.
 - d. Navigate to **Automation > Workflow Status** and track the status of the deployment request in the workflow status log until it completes successfully. It should take approximately 10 to 15 minutes to complete.

This procedure will result in a new 5G volume created on the SAN Volume Controller storage pool **mdiskgrp0** that is assigned to your windows host and mounted and formatted as the drive letter **H** within a few minutes with consistent results and no user errors.

For more information about workflows, see the following publications:

- *IBM Tivoli Storage Productivity Center Workflow User's Guide*
- Redpapers:
 - *Automated Storage Provisioning for Windows using Tivoli Storage Productivity Center and Tivoli Provisioning Manager*
 - *An Introduction to Storage Provisioning with Tivoli Provisioning Manager and Tivoli Storage Productivity Center*
- *White Paper: Tivoli Provisioning Manager and Tivoli Storage Productivity Center - Storage Provisioning Configuration and Exercises*

Archive and delete files

This scenario shows you how to reclaim storage space by archiving and deleting files that meet criteria you define.

You are running out of available storage capacity and need to find stale data which can be archived and deleted in order to reclaim space and make it available to more pertinent data. Specifically, you are **TPCUser** and you want to find files on the **C:** file system of **computer1** that are larger than **5MB**, and have not been accessed in **365** days.

Steps include the following tasks:

- Running a probe (if needed)
- Defining the constraint
- Running a scan
- Generating a constraint violation report
- Running an archive/backup on selected files in the report

Before you can use the IBM Tivoli Storage Manager archive or backup functionality you must ensure the prerequisites listed in “Requesting an archive or backup” on page 477 have been met.

To create an archive job, complete the following steps:

1. Run a Probe

You may not need to complete this step if a successful probe has been recently completed. Check the most recent day and time a probe was run successfully at **IBM Tivoli Storage Productivity Center > Monitoring > Probes**. To do this, expand the probe name to see the date the last successful probe was run. In most cases the default probe is sufficient.

If you do need to run a probe, complete the following steps:

- a. Expand **IBM Tivoli Storage Productivity Center > Monitoring > Probes**, and select **Create Probe**. The **Create Probe** window opens.
- b. Enter a description for the probe.
- c. In the **What to Probe** tab, select the resources you want to probe, in this case **Computer Groups > TPCUser.Default Computer Group > computer1**.
- d. In the **When to Run** tab, accept the defaults: **Run Now** and **Local time in each time zone**.

- e. In the **Alerts** tab, accept the default. If the probe fails, a notification will be sent to the error log automatically.
- f. Click **Save** to save the probe. Enter **Archive Backup Probe** as the probe name and click **OK**. An entry for the probe appears under **IBM Tivoli Storage Productivity Center > Monitoring > Probes > TPCUser.Archive Backup Probe**.
- g. The probe will begin running immediately. The number of times this job has been run and the date and time will be listed under the probe name. To see the status of the job, right-click the probe name and select **Update Job Status**.

2. Define Constraints

Define the criteria you are looking for with a constraint.

- a. Expand **Data Manager > Policy Management > Constraints**, right-click and select **Create Constraint**.
- b. In the **Description** field, enter a description of this constraint.
- c. In the **Filesystems** tab **Available** column, select **Filesystems > computer1 > C:** and move it to **Current Selections**.
- d. In the **File Types** tab, check **Forbid files matching these patterns**. All file types should be in the **Choose a file type** column. This will apply the constraint to all file types.
- e. In the **Users** tab, Check the **Forbid files owned by selected users**. All users should be listed in the **Available Users** column. This will apply the constraint to all users.
- f. In the **Options** tab, leave the default, **200**, in the text box in the **Violating File Limits** section. In the **Alerts on Files** section, select **number of days since last access greater than** and type **365**. Select **bigger than** and enter **5 MB**. Below the **Alerts on Files** section, **File Filter Text:** will list the values you picked in the **File Types**, **Users**, and **Options** tabs.
- g. In the **Alert** tab under **Triggering Conditions**, the only triggering condition is **Violating files consume more than**. Type **5** in the **Value** field. Select **Megabytes** from the **Value Units:** pull-down. Accept the default in **Triggered Actions**, which will send a message to the error log if the constraints you have defined are violated.
- h. Click the **Save** icon. Name the constraint **Archive File Constraint**. The constraint will be saved as **Data Manager > Policy Management > Constraints > TPCUser.Archive File Constraint**.

3. Run a Scan

Scan with the constraint violation you have created.

- a. Expand **Data Manager > Monitoring**, right-click **Scans** and select **Create Scan**. The **Create Scan** window opens.
- b. In the **Filesystems** tab, select **Filesystems > computer1 > C:/** from the **Available** column and move it to the **Current Selections** column.
- c. Do not select anything in the **Directory Groups** tab.
- d. In the **Profiles** tab, select all the profiles listed in the **Available Profiles** column and move them to both **Profiles to Apply to Filesystems** and **Profiles to Apply to Directories**.
- e. In the **When to Run** tab, select **Run Now** and **Local Time in each time zone**.
- f. In the **Alert** tab, **Scan Failed** is the only **Triggering Condition**. In the **Triggered Actions** area, select nothing to accept the default. If the job fails, a message will be sent to the log file automatically.

- g. Click the **Save** icon. Save the scan as **Archive Scan**. The scan is saved under **Data Manager > Monitoring > Scans > TPCUser.Archive Scan**. The scan starts to run immediately. Scans can take some time to complete depending on the number of storage assets scanned and the profiles used.
4. **Generate a Constraint Violation Report**
- Generate a report listing files that met your constraint violation criteria.
- a. Expand **Data Manager > Reporting > Usage Violations > Constraint Violations** and click **By Constraint**. The **Constraint Violations: By Constraint** window opens to the **Selection** tab.
 - b. In the **Report Filter Specifications** area, type **100** in the **Return maximum of** text field.
 - c. At the top right of the **Report Filter Specifications** area, click **Selection**. The **Select Resources** pop-up opens listing the available constraints.
 - d. Select **Archive File Constraint**. (You may have to de-select the other constraints.) Click **OK**.
 - e. Select the columns you want to appear in the report. By default, all columns will be listed in **Included Columns**. If you want to remove columns from the report, select them and move them to **Available Columns**.
 - f. Click **Generate Report**. The Constraint Violation Report is displayed on the monitor.
 - g. The page will stay the same with the addition of a **Constraints** tab which shows the **Constraint Violations: By Constraint** report. Click the magnifying glass to drill down to the **Constraint Violations: By Filesystem** report. Click the magnifying glass again and the **Largest Files: By Filesystem** report opens. This report, the **Constraint Violation Report**, will be used to create an Archive/Backup job.
5. **Run Archive/Backup on selected files in the report**
- Archive and backup files are automatically sent to the Tivoli Storage Manager server and are stored on that machine. Files are put into storage groups that are controlled through a Tivoli Storage Manager policy that the Tivoli Storage Manager server administrator usually creates.
- a. Select the files in the **Constraint Violation Report** that you want to archive or backup.
 - b. Right-click the selected files and select **Create a new archive\backup job**. The **Create New Archive/Backup Job** window opens.
 - c. In the **Options** tab, choose **Archive**. In this scenario the purpose is to archive and delete files that violate the constraints, so below the **Archive** selection, select **Delete After Successful Archival**. If, for some reason, you did **not** want to delete the files after they have been archived, do not select **Delete After Successful Archival**. If the archive is not successful, the files will not be deleted.
 - d. The files you have selected for archive/backup are shown in the **Files** tab.
 - e. In the **When to Run** tab select **Run Now**. In the **How to handle time zones** section, select **Local time in each time zone**.
 - f. In the **Alert** tab, in the **Triggering-Condition** area, the only triggering condition is **File Archive/Backup Failed**. Do not select anything in the **Triggered Actions** area to accept the default (a message sent to the error log).
 - g. Click the **Save** icon. Enter the name **Archive Log** and click **OK**. The job log is saved at **Data Manager > Monitoring > Policy Management > Archive/Backup > TPCUser.Archive Log**.

- h. A notice telling you that the job has been submitted appears. Click **OK**.
- i. Expand the job name to see the status of this particular job. You can judge the status by the color coding. When the job has finished, click the date and time to view the job log file.

Setting an archive/backup from the constraint

You can call archive backup when you set the constraint. However, if you set the delete option, it will automatically archive and delete **all** files that match the constraint criteria, which you probably do not want to do. At the very least you should not set archive/backup in the constraint the first time you run it.

To call for an archive/backup within a constraint, select **Archive/Backup** in the **Triggered Actions** area of the **Alert** tab in the **Create Constraint** window. Click **Define**. An **Archive/Backup Options** pop-up opens. Select **Delete After Successful Archival**. Go to the next tab and continue defining your constraint.

Related concepts:

“Requesting an archive or backup” on page 477

Request an IBM Tivoli Storage Manager archive and backup of the largest violating files identified by a constraint. Tivoli Storage Manager protects your data from hardware failures and other errors by storing backup and archive copies of data on auxiliary storage.

Related tasks:

“Creating probes” on page 336

Create and schedule a probe to job to gather information about your storage resources. You can define any number of probes to gather information about different storage resources at different times.

“Creating a file constraint” on page 471

Learn how to create a constraint that defines the acceptable and unacceptable file types, file sizes, and file owners for a computer or a set of computers in your environment.

“Defining an archive or backup in a constraint” on page 477

Learn how to request an IBM Tivoli Storage Manager archive and backup of the largest violating files identified by a constraint. Request a Tivoli Storage Manager archive and backup on the alert page of a constraint definition.

“Creating scans” on page 344

Create and schedule a scan to gather usage information about your storage resources. You can specify which file systems, computers, clusters, directories, databases, and other storage resources that you want to scan. You can define any number of scans to gather information about different storage resources at different times.

Create a profile with a filter

This scenario scans storage on a selected drive looking for the 100 largest media files, and summarizes space used by filesystem or directory and by owner.

There are a number of media files on computer1 that you suspect are taking up space that could be used for other purposes. You would like to search C:\ drive file systems and directories for the 100 largest files and create a report you can use to decide what action to take (contact the owner, archive and delete the files, and so forth). You will also publish this report to the web, where other people with permission levels greater to or equal to yours (TPCUser) can access it.

With this scenario, you perform the following tasks:

1. Create a profile to look for the largest media files
2. Create a filter within the profile
3. Create a new scan schedule using the profile
4. Create and run a batch report using the information gathered from the scan schedule
5. Save this report as an HTML file

For this scenario, complete the following steps:

1. **Create a new profile**
 - a. Expand **Data Manager > Monitoring**, right-click **Profiles** and select **Create Profile**.
 - b. The **Create Profiles** window opens. Enter a description for the profile.
 - c. In the **Statistics** tab, do the following:
 - 1) Under **Summarize space usage by**, select **filesystem or directory and owner**.
 - 2) Under **Accumulate History**, select **Per Scan** and type **10** as the number of days to keep.
 - 3) Under **Gather statistics by length of time since**, select **creation**.
 - 4) Under **Gather the information on the**, select **largest files** and type **100** in the text field.
 - d. In the **File Filters** tab, create the filter by doing the following:
 - 1) Highlight **(All Files Selected)** and click **New Condition**. The **Create Condition** window opens.
 - 2) Select **NAME** and **Matches any of** from the pull-down lists.
 - 3) Select **In any Directory**.
 - 4) In the **match** field, enter the types of files you are looking for in either a comma-separated list or by pressing return after each entry. As you press return, the file type will be listed in the box underneath the text field. For this scenario the file types are ***.avi**, ***.au**, ***.mpeg3**, and ***.gif**. Click **OK**. The **Create Profile** window re-opens, reflecting your choices.
 - e. Click the **Save** icon. Name the profile **Media Files Profile** and click **OK**. The profile will be saved as **Data Manager > Monitoring > Profiles > TPCUser.Media Files Profile**.
2. **Create the scan**
 - a. Expand **Data Manager > Monitoring**, right-click **Scans** and select **Create Scan**.
 - b. The **Create Scans** window opens. Type a description of the scan.
 - c. In the **Filesystems** tab, select **C:/** as the file system to scan.
 - d. Do not select anything in the **Directory Groups** tab.
 - e. In the **Profiles** tab, select **TPCUser.Media Files Profile** from Available Profiles and move it to **Profiles to apply to Filesystems**.
 - f. In the **When to run** tab, choose **Run Now** and select **Local time in each time zone**.
 - g. In the **Alerts** tab, the only **Triggering Condition** is **Scan Fails**. In the **Triggered Actions** area, select nothing to accept the default. If the job fails, a notification will be sent to the error log automatically.
 - h. Click the **Save** icon. Name the scan **Media Files Scan** and click **OK**. The scan will be saved as **Data Manager > Monitoring > Scans >**

TPCUser.Media Files Scan and the job will begin running immediately. This scan could take approximately 20 minutes to complete.

3. Create the batch schedule

Once you have run a scan with the new profile, the data is in the database but you have to go get it with a custom-designed report or batch schedule.

- a. Expand **IBM Tivoli Storage Productivity Center > My Reports**, right-click **Batch Reports** and select **Create Batch Report**.
 - b. The **Create Batch Report** window opens. Enter a description for the report.
 - c. In the **Available** column in the **Report** tab, select **Usage > Files > Largest Files > by Computer** and move that to **Current Selections**.
 - d. In the **Selection** tab, select **TPCUser.Media Files Profile** from the available profiles. By default, all columns are included in **Included Columns**. If you do not want a column in your report, select it and move it to **Available Columns**.
 - e. In the **Options** tab, make the following selections:
 - Under **Agent Computer Specification**, select **computer1**. This is the computer where the report will be run and the data stored. In most cases this is the Agent running on the Tivoli Storage Productivity Center server. If **computer1** is a Data Agent computer, you must accept the default report-destination path. If **computer1** is a V4.2 Storage Resource Agent or the Tivoli Storage Productivity Center server, you can accept the default path or you can specify a different path in the **Report destination** field. The default destination path varies depending on whether the computer is a Storage Resource Agent or the server.
 - Under **Report Type Specification**, select **HTML File**.
 - Select nothing under **Script**.
 - Under **Output File Specification**, enter the mask that will determine the output file name. It is recommended that you name the report something other than the default so that it can be easily identified by you and others. **{Report name}.{Report run date}** is recommended.
 - f. In the **When to run** tab, choose **Run Now** and select **Local time in each time zone**.
 - g. In the **Alerts** tab, the only **Triggering Condition** is **Report Failed**. In the **Triggered Actions** area, select nothing to accept the default. If the job fails, a notification will be sent to the error log automatically.
 - h. Click the **Save** icon. Name the batch report **Largest Media Files** and click **OK**. The batch report criteria will be saved under **IBM Tivoli Storage Productivity Center > Reports > Batch Reports > TPCUser.Largest Media Files**. The job will begin running immediately. The job number, date, and time will be listed under **TPCUser.Largest Media Files**.
4. Batch reports can be saved to any computer running Tivoli Storage Productivity Center. Reports are saved to a default location for Data Agents. If the computer is a V4.2 Storage Resource Agent or the Tivoli Storage Productivity Center server, you can accept the default path or you can specify a different path when you select batch report options.

Related tasks:

“Creating a profile” on page 354

Create a profile for use within a scan.

“Creating scans” on page 344

Create and schedule a scan to gather usage information about your storage resources. You can specify which file systems, computers, clusters, directories, databases, and other storage resources that you want to scan. You can define any number of scans to gather information about different storage resources at different times.

“Creating batch reports” on page 1021

Create a batch report.

“Selecting batch report options” on page 1022

Define the agent computer that you want to run the batch report on and specify the format of the file to be saved.

Creating a performance threshold within an alert

You can monitor the performance of your enterprise by creating alerts with performance thresholds for storage subsystems and SAN switches. By defining alerts that are triggered by performance thresholds, you can get timely information about performance issues in your enterprise.

To properly monitor the performance of your storage subsystem you will need to create the appropriate data collection jobs and alert definitions. The performance monitors collect data from the devices they are monitoring and at each sample interval the data is compared to the performance thresholds you have set. If the thresholds are exceeded, a notification is generated and the actions you defined in an alert are performed.

Note: The term Alert can be used one of three ways. An IBM Tivoli Storage Productivity Center *Alert definition* is defined at [Data | Disk | Tape] Manager > Alerting > [Type of Alert]. Within the Alert definition there is an *Alert tab*. Within the Alert tab there is a *Triggering Condition* area which gives the condition that will result in an *Alert message*. Most Triggering Condition areas in the Alert tab have only one Triggering Condition, such as **Probe Failed**. However, in Disk Manager some of the triggering conditions have thresholds, which are condition boundaries you set. For example, **Total I/O Rate Threshold** requires two upper and two lower boundaries or values. Only if these boundaries are matched or exceeded will an Alert message be generated and sent by the method you chose in *Triggered Actions*.

In this scenario, you are defining a performance threshold alert to monitor a disk subsystem (DS8000-21-7-1300271-IBM), for the **Total I/O Rate Threshold** condition, and generating, viewing, and printing the results.

To complete this scenario, you need to perform the following tasks:

1. Create a storage subsystem performance monitor
2. Create a storage subsystem alert definition
3. Generate and view a report showing the collected storage subsystem data

Note: You must run probes against storage subsystems prior to successful performance monitoring of those subsystems.

1. **Create a subsystem performance monitor**

- a. Expand **Disk Manager > Monitoring**, right-click **Subsystem Performance Monitor** and select **Create Subsystem Performance Monitor**. The **Create Storage Subsystem Monitor** page opens to the **Storage Subsystem** tab.
- b. Select the disk you want to monitor from the **Available subsystems** column and move it to the **Selected subsystems** column. Although we do not do it here, you can select more than one subsystem at a time.
- c. Click the **Sampling and Scheduling** tab. In the **Sampling** area, set the interval for gathering data to **1 hour**. Select **Continuing gathering data for** and select a value of **1 month**. In the **Scheduling** area, select **Begin immediately**.

Note: You do not want to sample too frequently as it increases the size of the database repository significantly. Monitors should be set to run for long periods, such as weeks or indefinitely.

- d. Click the **Alert** tab. **Monitor Failed** is the only **Triggering Condition** available. In **Triggered Actions**, select nothing to accept the default. If the job fails, a message will be sent to the log file automatically.
 - e. Click the **Save** icon. Name the monitor **SS Monitor**. It will be listed at **Disk Manager > Monitoring > Subsystem Performance Monitors > TPCUser.SS Monitor**. The job will begin running immediately.
2. **Create a storage subsystem performance alert**

- a. Expand **Disk Manager > Alerting**, right-click **Storage Subsystem Alerts** and select **Create Storage Subsystem Alert**. The **Create Storage Subsystem Alert** page opens.
- b. Click the **Alert** tab. Select **Total IO Rate Threshold** from the **Condition** drop down list in the **Triggering-Condition** area. Text fields will open for **Critical Stress**, **Warning Stress**, **Warning Idle**, and **Critical Idle**.

Note: The threshold values given in this scenario are examples and are **not** recommended values. They are for the purpose of illustration only. Your values will undoubtedly be different. Performance behavior varies tremendously between different devices and even different applications for different devices. We recommend monitoring your performance for several weeks and use this historical data to determine reasonable values for each threshold.

- In the **Critical Stress** field, type **3000** as the upper boundary for critical stress.
 - In the **Warning Stress** field, type **2200** as the upper boundary for warning stress.
 - Leave the **Warning Idle** field, the lower boundary for warning idle, blank.
 - Leave the **Critical Idle** field, the lower boundary for critical idle, blank.
- c. In the **Triggered-Actions** area, select **Email** to send an email to a mail address you specify.
 - 1) When you select **Email** the **Email Recipients** text box becomes active.
 - 2) Click **Add** and type an email address in the text box. You can type multiple email addresses.
 - 3) Click **Edit e-mail** to customize the email message. Note that errors sent to the error log are very generic, so you might want to be more descriptive in the text of the email.

- d. In the **Storage Subsystems** tab, select **Storage Subsystems > DS8000-21-7-1300271-IBM** in the **Available** column and move it to the **Current Selections** column.
 - e. To save the alert, click the **Save** icon. Name the alert **SS Alert**. It will be listed at **Disk Manager > Alerting > Storage Subsystem Alerts > TPCUser.SS Alert**. The conditions defined in the alert will be checked when the next probe runs against the subsystem.
3. **Generate and view a report**
- a. Expand **Disk Manager > Reporting > Storage Subsystem Performance**, and click **By Storage Subsystem**. The **Storage Subsystem Performance: By Storage Subsystem** window opens.
 - b. In the **Selection** tab, all performance metrics fields are selected by default. To remove a field from the report, click the metric name in **Included Columns** and move it to **Available Columns**. In the **Report Filter Specifications** area enter the number of storage subsystems upon which you want a report in the **Return maximum of** text field.
 - c. Click the **Generate Report** button. A **Storage Subsystem** tab will open containing the report. There will be one line in the report for each storage subsystem.
 - d. To save the report definition, click the **Save** icon. The **Save As** window opens. Name the report **SS Report** and click **OK**. The report will be saved at **IBM Tivoli Storage Productivity Center > My Reports > TPCUser's Reports > TPCUser.SS Report**.
 - e. Click the **Print** icon to print the current page on the local machine. You can only print one screen at a time.

Note: You might want to monitor the **Total IO Rate Threshold** condition for the storage subsystem by using constraints instead of thresholds. The Triggering Condition for a constraint would be **Violating Files Consume More Than**. This will alert you if any file goes over the condition value you set. In the **Value** field, type **2200**, which is the value of Warning Stress. In **Value Units**, select **Megabytes** from the pull-down list. See “Archive and delete files” on page 144 for more information on creating constraints.

You can use the above steps, with some minor changes, to create performance monitors and alerts for certain types of switches that have been discovered by Tivoli Storage Productivity Center. Switch performance monitors, for example, are created in the Fabric Manager node:

- To create a switch performance monitor, expand **Fabric Manager > Monitoring > Switch Performance Monitors**
- To create a switch performance alert, expand **Fabric Manager > Alerting > Switch Performance Alerts**.

Related tasks:

“Creating storage subsystem performance monitor” on page 377

Create a subsystem performance monitor to gather performance data for storage systems.

“Creating storage system alerts” on page 390

You can use storage systems alerts to monitor systems for changes in configuration or status. When the condition is detected during a probe, an alert is generated.

These alerts can be applied to storage subsystem groups and storage subsystems.

Related reference:

“Disk Manager reports” on page 961

Use Disk Manager reports to view the relationships among the components of a storage subsystem and the storage capacity for a computer, filesystem, storage subsystem, volume, Array Sites, Ranks, and disk.

Determine application performance problems

This scenario provides a method for finding out why an application has stopped performing as expected.

A previously normal application has stopped performing well. You must find out what is impacting the application so that you can take steps to correct the situation. You suspect a storage system problem.

In this scenario, you perform the following tasks:

- Probe Storage Resource agents and Data agents
- Probe the Storage Systems and start the Performance Monitor running
- Create a **Storage Subsystem: By Filesystem/Logical Volume** report
- Create a **Storage Subsystem Performance: By Volume** report
- Create a **Storage Subsystem Performance: By Array** report

Tip: Ordinarily, you run probes and performance monitors on both storage systems and switches, because the problem might be in either place. However, in this scenario you are only gathering information about storage systems.

Prerequisites:

- The Storage Resource agents and Data agents must be installed on the monitored computers before you can complete the tasks in this scenario. These are usually installed as part of the IBM Tivoli Storage Productivity Center installation process.
- Storage Resource agents, Data agents, and Arrays must have already been probed.
- Storage systems such as IBM System Storage DS8000 Storage Manager, IBM System Storage SAN Volume Controller, and the XIV systems must be connected through native interfaces. For other storage systems, CIM object manager (CIMOM) connections must be added and discovered.

Substitutable values that are used in this scenario:

- Device path = M:\ (from application owner or server administrator)
- Computer = ODCCL163.odclass.ibm.com (from Filesystem/Logical Volume report)
- Volume name = testvol-100 (id.1000)(from Filesystem/Logical Volume report)
- Storage system name = DS6000-1750-13AB44A-IBM (from Filesystem/Logical Volume report)

- Array = 1750.13AB24A-1 and A-2 (from the By Volume report, drill down to By Array report)

To track down application performance problems, collect information by running probes, performance monitors, and reports. Then use that information to determine actual performance and use values.

To perform this scenario, complete the following steps:

1. When the problem is reported, get computer information, device path information, and the application name from the application owner or the server administrator.
2. Probe the Storage Resource agents and Data agents.
 - a. Expand **IBM Tivoli Storage Productivity Center > Monitoring**. Right-click **Probes** and select **Create Probes**.
 - b. In the description field, type Probe for agents.
 - c. In the **What to Probe** page, expand 'Computers' in the **Available** box, select All Computers, and move that to **Current Selections**. This includes all Storage Resource agents and Data agents in the probe.
 - d. Click the **When to Run** tab and select **Run Repeatedly**, and set a regular time the probe is run. The following settings are up data collection for 2:00 AM every morning.
 - 1) Set the day in **Beginning at** a day ahead of the current day. If the day is February 21st, 2008, enter February 22, 2008, at 2:00 AM.
 - 2) Select **Repeat Every** and set **23 Hour(s)** to specify how often the probe is run. This setting allows time for the probe to shut down and start again.

In the **How to handle time zones** area, select **Use the timezone that the server runs in**, which uses the time zone where the Data server resides.
 - e. On the **Alert** page, **Probe Failed** is the only triggering condition that triggers an alert. In the **Triggered Actions** area, select **Windows Event Log** and set the **Event Type** to **Warning**.
 - f. Click the **Save** icon. Type AgentsProbe_1 as the probe name and click the **Save** button. The probe is listed in the navigation tree as **IBM Tivoli Storage Productivity Center > Monitoring > Probes > Administrator.AgentsProbe_1**.
 - g. To begin collecting data immediately, expand **IBM Tivoli Storage Productivity Center > Monitoring > Probes**, right-click **Administrator.AgentsProbe_1** and select **Run Now**. The status line for the probe is listed under **Administrator.AgentsProbe_1** by date and time.
A probe of the agents typically takes under a minute.
3. **Probe the Storage Systems**. Creating the Storage System probe is similar to creating the probe for the Storage Resource agents and Data Agents.
 - a. Expand **IBM Tivoli Storage Productivity Center > Monitoring**, right-click **Probes** and select **Create Probe**. In the **Description** field, type Probe for Storage Systems.
 - b. In the **What to Probe** tab Available area, expand **Storage Subsystems**. Select **All Storage Subsystems** and move it to Current Selections.
 - c. In the **When to Run** page, select **Run Repeatedly** beginning a day later than the current day at 3:00 AM.
 - d. Click the **Save** icon. Type StorageSystemsProbe_1 as the probe name and click the **Save** button. The listing for the probe is displayed in the

navigation tree as **IBM Tivoli Storage Productivity Center > Monitoring > Probes > Administrator.StorageSubsystemProbe_1**.

- e. To collect initial data, right click **Administrator.StorageSubsystemProbe_1** and select **Run Now**. Depending on the number and size of the storage systems to be probed, this probe might take up to 15 minutes.
4. **Begin running the storage system Performance Manager.** If you have a current Performance Monitor running and you have added another storage system, you must start a new Performance Monitor.
 - a. Expand **Disk Manager > Monitoring**, right-click **Subsystem Performance Monitor** and select **Create Performance Monitor**. In the **Description** field, type System performance monitor.
 - b. Click the **Storage Subsystems** tab. All storage systems that have had probes created and completed are listed under **Available subsystems**. Select all the available systems and move them to **Selected subsystems**.
 - c. Click the **Sampling and Scheduling** tab. In the Sampling area specify how frequently data is to be gathered and the length of the sampling window.
 - 1) For the Interval Length, type 5 in the minutes field.
 - 2) For Duration, select **Continue indefinitely**.
 - 3) Select **Begin Immediately** in the **Scheduling** area.
 - d. Click the **Alert** tab.
 - 1) In the **Condition** area, **Monitor Failed** is the only triggering condition.
 - 2) In the **Triggered Actions** area, select **Windows Event Log** with a **Warning** event type.
 - e. Click the **Save** icon and enter **SSPerformanceMonitor_1** as the performance monitor name. Click **OK**. The name of the storage system performance monitor is displayed at **Disk Manager > Monitoring > Subsystem Performance Monitor > Administrator.SSPerformanceMonitor_1**.

Tip: Setting performance alerts is not necessary for this scenario, but at this point you can review the historical performance data for your environment. Then you can define performance alerts with values that you consider out of the ordinary for your environment. You can set performance alerts for both storage systems and switches. After a performance alert is defined, future deviations will generate an alert and create an entry in the Performance Manager Constraint Violation report.

5. Gather information for at least two cycles (in this case, gather the information for at least 10 minutes). When information has been gathered, run the following reports to collect the data you can use to check performance.
 - **Disk Manager > Reporting > Storage Subsystems > Computer Views > By Filesystem/Logical Volume**
 - **Disk Manager > Reporting > Storage Subsystem Performance > By Volume**
 - **Disk Manager > Reporting > Storage Subsystem Performance > By Array**

The data from the reports is saved by date and time, but you might want to print copies of the reports for easier comparison.

- a. The **Storage Subsystem: By Filesystem/Logical Volume** report identifies the backend storage volumes associated with the file system. Variable values that are gathered from this report include values for computer, storage subsystem, and volume name.
 - 1) Run the report by expanding **Disk Manager > Reporting > Storage Subsystems > Computer Views** and click **By Filesystem/Logical Volume**.

- 2) Click the **Selection** tab, and in the **Report Filter Specifications** area, select **Volumes to Relate Filesystems/Logical Volumes to** in the list.
 - 3) Click the **Selection** button, clear the listed computers, and select your Device Path (M:\). Click **OK** to return to the **Selection** tab. When you run the report, the information is limited to your device path.
 - 4) By default all columns are listed as **Included Columns**. If there are columns you do not want in the report, select them and move them to **Available Columns**. Do **not** to move the "volume name" columns to **Available Columns**.
 - 5) Click **Generate Report**.
- b. The **Storage Subsystem Performance: By Volume** report shows you the actual performance metrics that are associated with the volumes the problem file system resides on.
- 1) Expand **Disk Manager > Reporting > Storage subsystem Performance** and click **By Volume**. By default all columns are listed as **Included Columns**. Do not move any of them to **Available Columns**.
 - 2) Click the **Selection** button. Click **Deselect All**, then find your system name (DS6000-1750-13AB24A-IBM) and component (the Volume Name - testvol-1000 (ID.1000)), and select that line. The information in the report is limited to that system and volume.
 - 3) Click **OK**. Click **Generate Report**.
- c. The **Storage Subsystem Performance: By Array** report shows the other volumes on the same array as the problem volume. It indicates whether the problem is specific to that one volume or if all of the volumes on the same array have a similar problem.

Note: Arrays are only supported on IBM TotalStorage Enterprise Storage Server, DS8000, and IBM System Storage DS6000systems.

- 1) Expand **Disk Manager > Reporting > Storage subsystem Performance** and click **By Array**.
 - 2) In the **Report Filter Specifications** area, select **Display latest performance data (by sample)**.
 - 3) Click the Selection button. Click **Deselect All** then find your arrays (Component) (1750.13AB45A-1 and A-2) and select those lines. The information in the report is limited to those arrays.
 - 4) By default all columns are listed as **Included Columns**. Do not move any of them to **Available Columns**.
 - 5) Click **Generate Report**.
6. Parts of the information that you use to determine the problem is in all of the reports. Look for the following things:
- Look at the I/O rate and the response time together for a complete picture of the performance. Ideally, you want a high I/O rate and low response time.
 - If the I/O rate is high but the response time is slowing over time, look at the cache hits to see how well cache is performing and, therefore, how often the I/O must go to disk.
 - If the I/O rate is low, look at the data rate. When there is a large number of I/O requests with high transfer sizes, the data rates become more important than I/O rates. In this case, look at the data rates and transfer sizes together.
 - Look at all the volumes in the same array to determine the I/O rate and response times. To determine these values, drill up from a selected volume in the 'By Volume' report to view the arrays that the volume is a part of. Then

drill down from the array to get a filtered 'By Volume' report of just the volumes on that array. Now compare the performance metrics against all the volumes in the same array. If all of them show the same performance problem, the performance bottleneck is the array. If only one volume has a performance problem, the performance bottleneck is somewhere else.

7. In this scenario, based on the I/O rates and the response times for the array over a time period, you determine that the array is being saturated with I/O and I/O should be off-loaded to other storage to improve performance.

Tip: You can also use Tivoli Storage Productivity Center to determine and project when more storage space is needed by monitoring I/O rates and response times over a time period.

Related information:

- *IBM Tivoli Storage Productivity Center User's Guide* or IBM Tivoli Storage Productivity Center Information Center
 - Configuring storage systems, see “Configure storage subsystems” on page 306.
 - Adding CIMOMs manually, see “Adding a CIM agent connection” on page 244.
 - Discovering CIMOMs, see “CIMOM discovery” on page 323.
 - Setting performance alerts, see “Creating alerts” on page 383.
- *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* or IBM Tivoli Storage Productivity Center Information Center
 - Installing agents, see Installing the agents.
 - Planning performance metrics, see Planning for performance metrics.

Monitor space on a remote server

Monitor a remote storage subsystem and take action when available space on a storage subsystem is less than 50GB.

A IBM Tivoli Storage Productivity Center server at a remote location (doremon) has storage on a DS4800 which is close to full. Monitoring has been set so that when the available storage space is less than 50GB an alert is generated to the server on which the Universal Agent is running. A system administrator can then use the Tivoli Enterprise Portal console to launch in context to the Tivoli Storage Productivity Center Wasted Space Report. Based on this report, orphan files and files that have not been accessed in a year on the DS4800 can be identified.

This scenario is designed for the Windows operating system. The following IBM products are used:

- **IBM Tivoli Monitoring** optimizes IT infrastructure performance and availability through a customizable workplace portal.
- **Tivoli Enterprise Portal** is the user interface for IBM Tivoli Monitoring. It is a portal-like application server that enables retrieval, manipulation, and analysis of data from the database server.
- A **Universal Agent** solution that reports Tivoli Storage Productivity Center asset information to IBM Tivoli Monitoring. This data is available for display in the Tivoli Enterprise Portal for reporting, charting and establishing situations in Tivoli Monitoring. A Universal Agent monitoring solution is provided by Tivoli Storage Productivity Center.
- The **launch in context** feature that allows you to pass parameters to the Tivoli Storage Productivity Center GUI. Based on the parameters you pass to the GUI,

you can automatically log into the specified Tivoli Storage Productivity Center server, navigate to a specific panel, and perform tasks on the specified server. launch in context is a feature of the Tivoli Storage Productivity Center user interface.

- **Tivoli Storage Productivity Center** provides a set of policy-driven automated tools for managing storage capacity, availability, events, performance and assets in your enterprise environment.

The following applications must be running:

- Tivoli Storage Productivity Center
- IBM Tivoli Monitoring, which includes Tivoli Enterprise Portal.
- The Tivoli Storage Productivity Center Universal Agent monitoring solution must be installed and configured per the readme. The readme is located in the TPCUA.zip file, which is on install disk1 in the tool folder.

Substitutable values used in this scenario:

- Tivoli Storage Productivity Center server = doremon
- doremon password = J24875
- IP address for doremon = 9.52.172.127
- User = Administrator
- <ITMHOME> = C:\ibm\ITM
- Situation name = DS4800_full
- Sampling interval = 15 minutes
- Wasted Space report name = WastedSpaceReport_Doremon

To complete this scenario, follow these steps:

1. Set up the IBM Tivoli Monitoring Universal Agent to monitor doremon.
This procedure is fully documented in the readme. Note that in real-life situations the Universal Agent would monitor much more than one server.
 - a. Start a command line session and change directories to <ITMHOME>.
 - b. Add doremon as a server to be monitored using the -config add command.
Enter the following command:

```
<ITMHOME>\tmaitm6\scripts\tpcua TPCUA.bat -config add -tpcDisplay  
"TPC Server Doremon" -tpcDBIP 9.52.172.127 -tpcDBPort 50000  
-tpcDBSchema TPC -tpcDBName TPCDB -tpcDBUser Administrator -tpcDBPass  
J24875
```
 - c. Schedule information collection using the Windows Scheduler.
 - 1) Enter the following command:

```
<ITMHOME>\tmaitm6\scripts\tpcua\TPCUA.bat -collect
```
 - 2) Go to **Start > All Programs > Accessories > System Tools** and click **Scheduled Tasks**. Double-click **Add Scheduled Task**. This starts the Scheduled Task Wizard. Click **Next**.
 - 3) Select **Browse**. Go to <ITMHOME>\tmaitm6\scripts\tpcua and select TPCUA.bat. Click **Open**.
 - 4) In the text field, type **TPCUA**. Select the **Daily** radio button. Click **Next**.
 - 5) In the **Start time** field, accept the default (current time). Select the **Every Day** radio button. In the **Start date** field, accept the default (current date). Click **Next**.
 - 6) Enter **Administrator** as the user name. Enter and confirm the password. Click **Next**.

- 7) Verify your selections. Click **Finish**. You will set the "repeat" time in the next steps.
 - 8) Go to **Start > All Programs > Accessories > System Tools > Scheduled Tasks**, right-click **Add Scheduled Task** and select **Properties**.
 - 9) In the "Schedule" tab, click **Advanced**. Select the **Repeat task** check box and specify 10 minutes as the time to repeat the collection. Click **OK**.
 - 10) Click **OK**.
- d. Import the metafile. On the command line, go to <ITMHOME>\tmai\itm6 and enter the following command:
- ```
kumpcon import TPC_Network.mdl
```

Once the metadata file has been imported, the Universal Agent collects Tivoli Storage Productivity Center data and provides it to IBM Tivoli Monitoring. It may be necessary to restart the Universal Agent after importing the .mdl file.

2. From IBM Tivoli Monitoring, set a situation for doremon that will alert you when the available storage space on the DS4800 falls below 50GB.
  - a. Using the Tivoli Enterprise Portal, open the Navigator to the Physical view. Expand **Enterprise > Windows Systems > DOREMON > Universal Agent > doremon:TPC\_NETWORK00**. and right-click **Select Situations**.
  - b. Select **STORAGE**. Right-click **STORAGE** and select **Situations**. The "Situations for - Storage" window will open.
  - c. Right-click **STORAGE** in the navigation tree and select **Create New**. The "Create Situation" window opens for you to enter the situation name and description.
  - d. Enter DS4800\_full for the name and "DS4800 storage subsystem nearly full on doremon" as the description. Click **OK**.
  - e. The "Select Condition" window opens.
    - 1) Select **Attribute Comparison** as the **Condition Type**
    - 2) Choose **STORAGE** in the Attribute Group area
    - 3) Choose **Available Capacity** in the Attribute Item area
    - 4) Click **OK**.
  - f. The "Situations for - STORAGE" window reopens to the **Formula** tab. The description for DS4800\_full may or may not be listed in the **Description** text box. If it is not, re-type it.
    - 1) In the Formula area, **Available Capacity** is listed with three cells below it. Click in the first cell. A formula area opens. Select a value of less than (<) 50.
    - 2) Leave **15 min** as the default **Sampling interval**.
    - 3) To play a sound to alert you, select **Enable** in the Sound area.
  - g. Click the **Action** tab.
    - 1) For **Action Selected**, select **System Command** and type "net send DS4800 managed by Doremon is nearly full" in the **System Command** text box.
    - 2) For **If the condition is true for more than one monitored item**, select **Only take action on the first item**.
    - 3) For **Where should the Action be executed (performed)**, select **Execute the Action at the Managed System (Agent)**. This will send the command window with message and the sound to the system on which the Universal Agent monitoring Doremon is running.

- 4) For **If the condition stays true over multiple intervals**: select **Do not take action twice in a row (wait until situation goes false then true again)**.
  - h. Click the **Distribution** tab. Verify that **doremon:TPC\_NETWORK00** is the correct assigned agent. If there is more than one assigned agent, move any incorrect agents to the Available Managed Systems area.
  - i. Click **OK**. The "Situations for - STORAGE" navigation tree will list **STORAGE > DS4800\_full**. If you should be returned to the "Welcome SYSADMIN" window, expand **Enterprise > Windows Systems > DOREMON > Universal Agent > doremon:TPC\_NETWORK00** and right-click **Situations**. The "Situations for - STORAGE" window will open and **DS4800\_full** will be listed.
3. IBM Tivoli Monitoring checks the data provided by the Universal Agent every 15 minutes. When storage on the DS4800 falls below 50GB, you will be notified by a sound and the message you entered in the Action tab in a command window on the system on which the Universal Agent monitoring Doremon is running.
  4. Run a Wasted Space report. When you are alerted that storage space is less than 50GB on the DS4800, use Tivoli Enterprise Portal to create a definition for launch in context that will launch a Wasted Space report.
    - a. In the Tivoli Enterprise Portal window, expand **Enterprise > Windows Systems > DOREMON > Windows OS**, right-click **Disk** and select **Launch**.
    - b. If you **have not** previously created a launch in context definition for the Wasted Space report task, complete the following steps:
      - 1) In the Tivoli Enterprise Portal window, expand **Enterprise > Windows Systems > doremon > Windows OS** Right-click **Disk** and select **Launch**.
      - 2) The "Create or Edit Launch Definitions" window opens. Click **Create New**.
      - 3) In the **Name** text field, type `WastedSpaceReport_Doremon`,
      - 4) For the **Target** text box, select **Browse** and go to `C:\Program Files\IBM\TPC\gui\tpc.bat`.
      - 5) In the **Arguments** text field, enter the parameters and values you will use to run the Wasted Space report:  
`-user admin -passwd J24875 -server doremon:5000 -function TPC.reports.data.wasted_space`
      - 6) For the **Start in** location, click **Browse** and select `C:\Program Files\IBM\TPC\gui`.
      - 7) Click **Apply** to save the report definition as **WastedSpaceReport\_Doremon** in the **Existing Launch Definitions** panel.
      - 8) Click **Launch** to run the report.
    - c. If you **have** previously created a definition for **WastedSpaceReport\_Doremon**, select the definition from the **Existing Launch Definitions** panel and click **Launch**.
  5. The Wasted Space report is displayed within the Tivoli Storage Productivity Center GUI, showing the totals for orphaned and stale files that can probably be removed to free up space on the DS4800.

The following documents were referenced for this scenario:

- A Universal Agent Solution readme

- IBM Tivoli Monitoring User's Guide or IBM Tivoli Storage Productivity Center Information Center:
  - Creating a situation
- IBM Tivoli Storage Productivity Center User's Guide or IBM Tivoli Storage Productivity Center Information Center:
  - Launch in context
  - Wasted Space report

## Identifying performance hot spots

This scenario uses the Storage Optimizer to analyze storage systems to determine if there are performance hot spots. This scenario also illustrates how to use the Storage Optimizer to get recommendations for improving storage system performance.

You want to analyze your two IBM System Storage DS8000 Storage Manager storage systems to determine if there are performance hot spots. You want to identify the hot spots and get recommendations on how to improve storage system performance.

### Prerequisites for using Storage Optimizer

- IBM Tivoli Storage Productivity Center Standard Edition license
- One or more of these supported systems: DS8000, IBM System Storage DS6000, IBM System Storage DS5000 Storage Manager, IBM System Storage DS4000, IBM TotalStorage Enterprise Storage Server, IBM System Storage SAN Volume Controller, or IBM Storwize V7000
- At least one week's worth of performance monitoring data gathered in advance of using Storage Optimizer. Providing a longer time interval for data collection increases the accuracy of the analysis and recommendations.

### Monitor the performance of your DS8000 systems

At least one week in advance, perform performance monitoring for your two DS8000 storage systems. The general steps are as follows:

1. Use the Configure Devices wizard to enter connection information for the native interfaces of the DS8000 systems. See “Configure storage subsystems” on page 306 for details of how to use the Configure Devices wizard. If you want to monitor SAN Volume Controller systems or the XIV systems, you can also use the Configure Devices wizard to enter connection information for their native interfaces. For systems that connect through CIM object manager (CIMOM) connections, run a discovery for storage system CIMOMs. See “Discover storage subsystem, tape library, and fabric information with CIM agents” on page 330 for more information.
2. Create a probe. See “Creating probes” on page 336 for more information.
3. Create a storage system performance monitor and let it run for at least one week. See “Creating storage subsystem performance monitor” on page 377 for more information.

To perform this scenario, complete the following steps:

#### Part 1: Create an analysis report that includes both DS8000 storage systems.

1. At least one week in advance, run a storage system performance monitor to collect performance data for both DS8000 systems.
2. Navigate to **IBM Tivoli Storage Productivity Center > Analytics**.

3. Right-click **Storage Optimizer** and select **Create Analysis**. The Create Analysis panel lets you specify the storage systems that you want to analyze.
4. Under Create Analysis, enter "Analyze DS8000 systems" in the Description field.
5. Click **Add** in the Select Storage pane. The Optimizer Selection panel opens. This provides a topology view of the current storage network.
6. Select both of your DS8000 storage systems.
7. Click **OK** to return to the Storage Optimizer Create Analysis panel with the DS8000 systems displayed in the Select Storage pane.
8. Set the Performance Time Interval to start one week prior to the current date.
9. For How often to run, select **Run Now**.
10. Choose **File > Save** to save and submit the analysis report job.
11. When prompted to enter a name for the report job, enter "Analyze DS8000 systems" and click **OK**.
12. View the analysis report job status for "Analyze DS8000 systems" under the **Storage Optimizer** node in the navigation tree.

### Part 2: View the analysis report and identify hot spots

When the analysis report job "Analyze DS8000 systems" is complete (status is green), select it to view the analysis report for your DS8000 systems. The heat maps display color-coded cells that represent the projected utilization of both DS8000 systems, as measured against the desired performance threshold of 80%. Each cell represents a storage pool.

1. From the list of components for **Heat Map based on**, choose Hard Disk to update the heat maps to display the analysis data for that component.
2. Move the performance threshold slider to see its affect on the heat maps.
3. Click the **Performance Tables** tab to display a tabular view of the data, if desired.
4. In the Hard Disk heat map for one of your DS8000 systems, you notice a red cell, or "hot spot". Move your cursor over that cell to display the details for the storage pool named Pool 1. The details for Pool 1 indicate that it is exceeding the desired performance threshold of 80%.

Just to the left of the red cell is a green cell for Pool 2. The green color indicates that Pool 2 is performing at less than or equal to 25% of the established performance threshold.

You think that you might be able to improve the performance of Pool 1 by moving some of its volumes to Pool 2, and you know that Pool 2 has some amount of available capacity.

### Part 3: Create an optimization report that includes recommendations on how to improve the performance of Pool 1.


1. In the analysis report, select the red cell which represents Pool 1.
2. Click >> next to **Source Entities** to move Pool 1 into the list of source entities.
3. Select the green cell which represents Pool 2.
4. Click >> next to **Target Entities** to move Pool 2 into the list of target entities.
5. Choose **File > Save** to save and submit the optimization report job.
6. When prompted to enter a name for the optimization report job, enter "Optimize DS8000 systems" and click **OK**.



7. Under **View Previously Run Optimization Reports**, click **Refresh Job Status** to update the report job status.

#### **Part 4: View the optimization report and recommendations.**

When the optimization report job "Optimize DS8000 systems" is complete (job

status is "Success"), click  to display the optimization report. The report includes a recommendation for how to improve the performance of Pool 1.

The heat maps in the optimization report display the performance utilization of Pool 1 and Pool 2 before and after implementing the recommendation. The heat maps show you that the red cell representing Pool 1 changes to green after implementing the recommendation. The heat maps also show you that the green cell representing Pool 2 changes to yellow after implementing the recommendation. This means that Pool 2 will be still be performing at less than or equal to 75% of the performance threshold after you implement the recommendation.

To view the effects on the before and after heat maps of implementing only some of the recommendations:

- Remove the check mark next to **Select all recommendations** and select a subset of the recommendations from the recommendations table.

To print the optimization report:

- Click the Print button in the optimization report.

The printed report includes additional information such as SVC sample scripts to help you implement the recommendations using the SAN Volume Controller command-line interface.

## **Retiring underutilized storage systems**

This scenario illustrates how you can use the Storage Optimizer to create recommendations about retiring underutilized storage systems and migrating their data to other systems in your data center.

Your data center has seven IBM System Storage DS8000 Storage Manager storage systems in use. Using various IBM Tivoli Storage Productivity Center capacity reports, you have determined that two of the systems are underutilized. You want to create a report that provides recommendations to help guide you in retiring the underutilized systems and migrating the data into your remaining five systems.

#### **Prerequisites for using Storage Optimizer**

- IBM Tivoli Storage Productivity Center Standard Edition license
- One or more of these supported IBM systems: DS8000, IBM System Storage DS6000, IBM System Storage DS5000 Storage Manager, IBM System Storage DS4000, IBM TotalStorage Enterprise Storage Server, IBM System Storage SAN Volume Controller, or IBM Storwize V7000
- At least one week's worth of performance monitoring data gathered in advance of using Storage Optimizer. Providing a longer time interval for data collection will increase the accuracy of the analysis and recommendations.

#### **Monitor the performance of your DS8000 systems**

At least one week in advance, perform performance monitoring for all your DS8000 storage systems. The general steps are as follows:

1. Use the Configure Devices wizard to enter connection information for the native interfaces of the DS8000 systems. See “Configure storage subsystems” on page 306 for details of how to use the Configure Devices wizard. If you want to monitor SAN Volume Controller systems or the XIV systems, you can also use the Configure Devices wizard to enter connection information for their native interfaces. For systems that connect through CIM object manager (CIMOM) connections, run a discovery for storage system CIMOMs. See “Discover storage subsystem, tape library, and fabric information with CIM agents” on page 330 for more information.
2. Create a probe. See “Creating probes” on page 336 for more information.
3. Create a storage system performance monitor and let it run for at least one week. See “Creating storage subsystem performance monitor” on page 377 for more information.

To perform this scenario, complete the following steps:

**Part 1: Create an analysis report for the seven DS8000 storage systems in your data center.**

1. At least one week in advance, run a storage system performance monitor to collect performance data for all DS8000 systems.
2. Navigate to **IBM Tivoli Storage Productivity Center > Analytics**.
3. Right-click **Storage Optimizer** and select **Create Analysis**. The Create Analysis panel lets you specify the storage systems that you want to analyze.
4. Under Create Analysis, enter "Analyze DS8000 systems" in the Description field.
5. Click **Add** in the Select Storage pane. The Optimizer Selection panel opens. This provides a topology view of the current storage network.
6. Select all your DS8000 storage systems.
7. Click **OK** to return to the Storage Optimizer Create Analysis panel with the DS8000 systems displayed in the Select Storage pane.
8. Set the Performance Time Interval to start one week prior to the current date.
9. For How often to run, select **Run Now**.
10. Choose **File > Save** to save and submit the analysis report job.
11. When prompted to enter a name for the report job, enter "Analyze DS8000 systems" and click **OK**.
12. View the analysis report job status for "Analyze DS8000 systems" under the Storage Optimizer node in the navigation tree.

**Part 2: View the analysis report.**

When the analysis report job "Analyze DS8000 systems" is complete (status is green), select it to view the analysis report for your DS8000 systems.


- For **Heat Map based on**, choose the **Space** option to view the heat maps in the analysis report based on disk space utilization. The heat maps confirm that the majority of pools in systems 6 and 7 are performing at less than 75% of capacity.

**Part 3: Create an optimization report that includes recommendations on how to retire the underutilized systems 6 and 7 and migrate the data to your remaining five systems.**

1. In the heat maps, select all pools for storage systems 6 and 7 using the **Select all pools for storage subsystem** checkbox.
2. Click >> next to **Source Entities** to move systems 6 and 7 into the list of source entities.
3. Click **Make retirement recommendations for the selected entity**.
4. In the heat maps, select all pools for storage systems 1 through 5 using the **Select all pools for storage subsystem** checkbox.
5. Click >> next to **Target Entities** to move systems 1 through 5 into the list of target entities.
6. Choose **File > Save** to save and submit the optimization report job.
7. When prompted to enter a name for the report job, enter "Retire DS8K" and click **OK**.
8. Under **View Previously Run Optimization Reports**, click **Refresh Job Status** to update the report job status.

#### Part 4: View the optimization report and recommendations.

When the optimization report job "Retire DS8K" is complete (job status is

"Success"), click  to display the optimization report. If it is possible to retire systems 6 and 7, the report will include recommendations for retiring systems 6 and 7 and migrating the data into systems 1 through 5.

- For **Heat Map based on**, choose **Space** to view the before and after optimization heat maps based on disk space utilization.
- To view the effects on the before and after heat maps of implementing only some of the recommendations, remove the check mark next to **Select all recommendations** and select a subset of the recommendations from the recommendations table.

The before and after heat maps illustrate the effects of retiring systems 6 and 7 and migrating the data into systems 1 through 5. The heat maps also confirm that even after retiring systems 6 and 7, the remaining systems will still perform within the limits of the desired performance threshold.

To print the optimization report:

- Click the Print button in the optimization report.

The printed report includes additional information such as SVC sample scripts to help you implement the recommendations using the SAN Volume Controller command-line interface.

## Tivoli Storage Productivity Center for Replication scenarios

These scenarios illustrate how to perform specific replication tasks using IBM Tivoli Storage Productivity Center for Replication.

The provided scenarios set out situations that can be solved using IBM Tivoli Storage Productivity Center. When using these scenarios you should remember the following:

- The scenarios are intended to be step-by-step instructions that you follow to complete a specific task. Each scenario contains multiple steps that must be followed in the order given.

- Specific values such as user names, passwords, and IP Address are given in the scenarios only for illustration purposes. You must replace them with values appropriate for your environment.
- The scenarios do not provide in-depth conceptual information about the tasks.

### **Identifying and investigating a suspended Metro Mirror session**

IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication work together to help meet your storage management needs and automate your replication management tasks. This scenario shows you how to use IBM Tivoli Storage Productivity Center to view and investigate a suspended Metro Mirror between two IBM DS8000 storage systems that was detected by IBM Tivoli Storage Productivity Center for Replication.

You want to use IBM Tivoli Storage Productivity Center to analyze alerts generated by events detected by IBM Tivoli Storage Productivity Center for Replication. You want to open a page in IBM Tivoli Storage Productivity Center for Replication from the IBM Tivoli Storage Productivity Center user interface.

#### **Prerequisites:**

- Two DS8000 storage systems.
- Use IBM Tivoli Storage Productivity Center for Replication to mirror a database application between the DS8000 storage systems.
- Set up IBM Tivoli Storage Productivity Center to monitor DS8000 storage systems. The general steps are as follows:
  1. Run a CIMOM discovery to identify the storage systems managed by CIMOMs in your environment. See information about discovering storage resources in the *IBM Tivoli Storage Productivity Center for Replication for System z User's Guide* for more information about how to discover storage systems.
  2. Create and schedule a probe to run against the storage systems. See information about collecting storage statistics (probes) in the *IBM Tivoli Storage Productivity Center for Replication for System z User's Guide* for more information about how to define probes.

IBM Tivoli Storage Productivity Center includes a predefined set of alerts that correspond to events detected by IBM Tivoli Storage Productivity Center for Replication. In this scenario, IBM Tivoli Storage Productivity Center for Replication detects that a Metro Mirror session just went into a suspended state. That event is communicated to IBM Tivoli Storage Productivity Center and recorded in the alert log. Use the alert log to view information about that alert and start the IBM Tivoli Storage Productivity Center for Replication user interface. Use IBM Tivoli Storage Productivity Center for Replication to determine that the paths between your DS8000 devices in a Metro Mirror session were removed and resolve the problem.

The following tasks are described in this scenario:

1. IBM Tivoli Storage Productivity Center: Set up notification methods for replication alerts.
2. IBM Tivoli Storage Productivity Center: Use the alert log to view information about a suspended Metro Mirror session that was detected by IBM Tivoli Storage Productivity Center for Replication.
3. IBM Tivoli Storage Productivity Center for Replication: Start IBM Tivoli Storage Productivity Center for Replication from the IBM Tivoli Storage Productivity Center user interface to access the Session Details panel for the alert and resolve the problem.

To use IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication to analyze storage events, complete the following steps:

1. IBM Tivoli Storage Productivity Center: Set up notification methods for replication alerts.
  - a. Start IBM Tivoli Storage Productivity Center.
  - b. In the navigation tree, expand **Replication Manager > Alerting** and select **Replication Alerts**.
  - c. Click **TPCUser.Replication Suspending Event**. The Edit Replication Alerts panel is displayed. Use this panel to determine how you are notified if IBM Tivoli Storage Productivity Center for Replication detects a suspending event in your storage environment.
  - d. In the **Triggered Actions** section, select the notification methods for the alert.

By default, all suspending events that are detected by IBM Tivoli Storage Productivity Center for Replication are recorded as entries in the Alert Log panel under **IBM Tivoli Storage Productivity Center > Alerting** . If you want to be notified using additional methods when the alert log is updated for a suspending event, select any of the following options:

- Select **SNMP Trap** to generate an SNMP trap message to any NMS, console, or terminal to indicate the triggering of a suspending alert.
- Select **TEC/OMNIBus Event** to send an alert to the Tivoli Event Console or OMNIBus.
- Select **Login Notification** to be notified of a suspending event when you log in to Tivoli Storage Productivity Center.
- Select **Windows Event Log** or **UNIX SYSLOG** to write the suspending alert to the operating system log file.
- Select **Run Script** to run a script when a suspending alert is triggered.
- Select **Email** to indicate the email addresses that receive notification when a suspending alert is triggered.

You are not required to select any of these options. Triggered alerts from IBM Tivoli Storage Productivity Center for Replication are automatically recorded in the alert log. These options are meant to provide you with additional methods for receiving alerts.

- e. Click the **Save** icon to save the alert definition. When IBM Tivoli Storage Productivity Center for Replication detects that a session has suspended, this alert is triggered, an entry is added to the IBM Tivoli Storage Productivity Center alert log, and any additional triggered actions are run.
2. IBM Tivoli Storage Productivity Center: Use the alert log to view information about a suspended Metro Mirror session.

This step occurs after IBM Tivoli Storage Productivity Center for Replication has detected a suspended session and triggered an alert in the IBM Tivoli Storage Productivity Center user interface.

- a. In the navigation tree, expand **IBM Tivoli Storage Productivity Center > Alerting** and select **Alerting Log**. A visual indicator is displayed next to each alert type under this node. In this scenario, a red, filled-in circle is displayed next to the **Replication** node to indicate that IBM Tivoli Storage Productivity Center has received an alert from IBM Tivoli Storage Productivity Center for Replication that is unresolved.

b. Click the **Replication** node. The content pane displays a table that contains a row for each alert received from IBM Tivoli Storage Productivity Center for Replication. The **Alert Type** column in this table indicates the type of alert that was triggered.

c. Click the magnifying glass icon next to the alert that displays **Suspending Event** in the **Alert Type** column. The Detail for Alert panel displays a description of the event that triggered the alert:

```
IWNR1958W [2009/03/04 15:11:49.250 EST] Session session_name has changed to the SEVERE status due to an unexpected error.
```

where *session\_name* represents the name of the session in IBM Tivoli Storage Productivity Center for Replication.

d. Click the back arrow in the tool bar to return to the Alert Log panel.

3. Start IBM Tivoli Storage Productivity Center for Replication from a suspended event alert in the alert log.

The Session Details page indicates that the session went suspended because the paths between the devices were dropped.

a. On the **IBM Tivoli Storage Productivity Center > Alerting > Alerting Log** page, right-click the row for that represents the **Suspending Event** alert type. The menu that is displayed contains the actions that you can perform against the alert log entry. One of the actions is to start the IBM Tivoli Storage Productivity Center for Replication user interface.

b. Click **Launch Replication Session Details** on the menu. The IBM Tivoli Storage Productivity Center for Replication user interface starts in a Web browser and displays the Session Details page. Use this page to view detailed information about the suspended event.

Single sign-on is an authentication process that enables you to enter one user ID and password to access multiple applications. Single sign-on integrates with the launch in context feature to enable you to move smoothly from the IBM Tivoli Storage Productivity Center user interface to a specific location in the IBM Tivoli Storage Productivity Center for Replication user interface without re-entering a user ID and password.

For example, this scenario shows how you can use single sign-on and launch in context to open a page in the IBM Tivoli Storage Productivity Center for Replication user interface based on a triggered alert in IBM Tivoli Storage Productivity Center.

c. Use IBM Tivoli Storage Productivity Center for Replication to resolve the problem that triggered the alert.

You can use steps similar to the ones above to use IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication together to investigate and resolve different conditions and events that occur within your storage environment.

## Performing planned maintenance scenarios

Scenarios for a planned outage consist of the removal of a unit from service to perform scheduled work on specific components for a predetermined duration.

### Planned outage of local (H1), moving production to site 2 (H2) and optional return:

In this scenario, your goal is to shut down the H1 box for maintenance, or for a planned power outage. As a result, you must move production to the H2 site, with the option of switching back to the H1 site after it is restored.

Follow these steps to carry out this scenario.

1. **Start H1->H2->H3**

Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. The session will reach a Prepared state so there is a consistent point on Sites 2 and 3.

**Note:** Since the outage is a planned outage, stop production I/O to H1.

2. **Suspend**

This command creates the planned outage by suspending the Metro Mirror relationship between H1 and H2. When this step is complete, the session state should be suspended and recoverable.

3. **Recover H2**

This command makes H2 target-available. When the session state is Target Available, the production I/O can then be moved to write to the H2 volumes.

4. **Start H2->H3**

This command starts a Global Mirror session during the outage of H1. Production I/O can continue on H2. After the session reaches the Prepared state, a consistent copy of the data resides on Site 3, and can be used in the event of a disaster.

5. **Start H2->H1->H3**

Issue this command when H1 is available and ready to be brought back into the configuration. This command starts a Metro Global Mirror session with H1 and H2 roles reversed from the original configuration. Production I/O can continue on H2. After the session reaches the Prepared state, both Site 1 and Site 3 will have a consistent copy of the data.

**Note:** The following steps are the steps you must follow after you decide to switch production I/O to write on Site 1 again. Stop production I/O on H2 before continuing to the next step.

6. **Suspend**

This command suspends the Metro Mirror relationship between H2 and H1. When this step is complete, the session state should be suspended and recoverable.

7. **Recover H1**

This command makes H1 target-available. When the session state is target-available at H1, run production I/O to H1.

8. **Start H1->H2->H3**

This command returns the session to its original configuration, copying data from Site 1, to Site 2, to Site 3.

**Planned outage on site 2 (H2) and return to three-site configuration:**

In this scenario, there is a planned outage expected on site 2. Disaster recovery capability is still needed so the copy will be setup between site 1 and site 3. After site 2 is brought back up, the session can then be brought back up to its original configuration.

Follow these steps to carry out this scenario:

1. **Start H1->H2->H3**

Run this command to begin Metro Global Mirror on your system when the production I/O is running on H1. After the session is in a Prepared state, both sites 2 and 3 will have a consistent copy of the data.

**Note:** Production I/O does not stop running to site 1 to perform the following steps.

2. **Start H1->H3**

This command starts a Global Mirror session between site 1 and site 3, during the outage of H2. Production I/O can continue on H1. After the session has reached a Prepared state, a consistent copy of the data will reside on Site 3.

**Note:**

- If running DS8000 devices on H1 with the correct level of microcode, and enough time has passed since the initial start of the session, the creation of the H1 to H3 pairs will *not* require a full copy.
- Data on Site 2 is no longer consistent at this point.
- After Site 2 is brought back up and is ready to hold a consistent copy of the data, go to the next step to return to the full three-site configuration. Production I/O never has to stop writing to H1 in this scenario.

3. **Start H1->H2->H3**

After the session reaches a Prepared state, site 2 and 3 will have a consistent copy of the data again.

**Planned outage of local (H1) with production move to remote (H3):**

In this scenario, there is a planned outage for Site 1 (H1). Site 2, however, does not have the resources to run production, so production will be moved to Site 3 during the outage at Site 1. When the outage is complete, production can then be moved back to Site 1 and resume the full three-site configuration.

Complete the following steps to carry out this scenario.

1. **Start H1->H2->H3**

Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. Once the session is in a Prepared state, both sites 2 and 3 will have a consistent copy of the data.

**Note:** Because this is a planned outage at Site 1, stop production I/O to Site 1 (H1) before continuing to the next step.

2. **Suspend**

This command creates the planned outage by suspending the Metro Mirror relationship between H1->H2. When this step is complete, the session state should be suspended and recoverable.

3. **Recover H3**

This command makes H3 Target Available. When the session state is Target Available at H3, production I/O can be moved to H3.

**Note:** When Site 1 is ready to resume production, continue with the following steps. Production I/O can continue to be written to H3 through the following step.

4. **Start H3->H1->H2**

Issue this command when H1 and H2 are available and ready to be brought back into the configuration. Changes from production at the remote site (H3) flow from the local site (H1) to the intermediate site (H2). The H3-H1 and H1-H2 pairs are Global Copy. The session will remain in a preparing state.



**Note:** While the session is Preparing, you can decrease the amount of time that production I/O has to be stopped by waiting for progress on the H3-H1 role pair to be at or near 100%. The next step will drive H3-H1 to sync, which depending on the percentage of the copy that is complete, can take a significant amount of time. When you are ready to switch back to Site 1, stop production I/O on H3 before continuing to the next step.

5. **Suspend**

This command suspends Global Copy relationship between H3->H1 and H1->H2. When this step is complete, the session state should be suspended and recoverable.

6. **Recover H1**

This command makes H1 Target Available. When the session state is Target Available at H1, production can be switched back to H1.

7. **Start H1->H2->H3**

This command will resume copy from Site 1 back to sites 2 and 3. Once the session reaches a Prepared state, there will be a consistent copy on Site 2 and Site 3.

**Planned outage of remote (H3) and returning:**

In this scenario, the remote site H3 has a planned outage, or consistency groups from the Global Mirror session need to be temporarily paused.

Complete the following steps to carry out this scenario.

1. **Start H1->H2->H3**

Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. After the session is in a Prepared state, both sites 2 and 3 will have a consistent copy of the data.

**Note:** Production I/O to H1 can continue during the following steps.

2. **Suspend H2->H3**

Issue this command to suspend only the Global Mirror portion of the session. The Metro Mirror continues between H1 and H2. This causes the session state to be put into a SuspendedH2H3 state.

3. **Start H1->H2->H3**

This command restores the original configuration by restarting the Global Mirror portion of the session.

**Recovering from a disaster scenarios**

Scenarios for disaster recovery (or unplanned outages) describe the steps to recover from the loss of a complete site and restart operations after a disaster or disruption.

**Unplanned outage of local (H1), production move to intermediate (H2):**

In this scenario, there was an unplanned outage of the local site (H1), and production must be moved to Site 2 (H2) until Site 1 can be brought back up.

Complete the following steps to carry out this scenario.

**Note:** To ensure consistency in the event of a rolling disaster where the IBM Tivoli Storage Productivity Center for Replication server is lost on Site 1 before freezing all the copies in the session, enable the Metro Mirror heartbeat function.

1. **Start H1->H2->H3**

Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. Once the session is in a Prepared state, both sites 2 and 3 will have a consistent copy of the data.

2. **Unplanned outage of H1**

This event causes an unplanned outage of H1, which causes the session to go suspended. It is recoverable if the session was in a Prepared state when the outage occurred. If the state was Preparing when the outage occurred, the session might be recoverable at Site 2 but will not be recoverable at Site 3.

3. **RecoverH2**

This command makes H2 Target Available. When the session state is Target Available at H2, production can be switched to H2.

4. **Start H2->H3**

This command starts a Global Mirror session from Site 2 to Site 3 during the outage of H1. Once the session reaches a Prepared state, consistent data should reside on Site 3.

**Note:** Do not perform the following step until Site 1 is back up and ready to receive a copy of the data that has been written to Site 2 during the outage.

5. **Start H2->H1->H3**

This command starts an MGM session with the H1 and H2 roles reversed from the original configuration. Production I/O can continue on H2. Wait for the session to go to the Prepared state, at which point there will be a consistent copy of the data on Site 1 and Site 3.

**Note:** The following steps can be used to return to the original configuration so that production I/O can be switched back to Site 1. Stop production I/O on H2 before continuing to the next step.

6. **Suspend**

This command suspends the Metro Mirror relationship between H2 and H1. When this step is complete, the session state should be suspended and recoverable.

7. **RecoverH1**

This command makes H1 Target Available. When the session state is Target Available at H1, production I/O can be switched back to H1.

8. **Start H1->H2->H3**

This command restores the original configuration. Once the session reaches a Prepared state, a consistent copy of the data will reside on Site 2 and Site 3.

**Unplanned outage of Site 2 (H2):**

In this scenario, an unplanned outage of Site 2 (H2) occurs. Disaster Recovery capability is still required, so a Global Mirror session must be set up between Site 1 (H1) and Site 3 (H3) while Site 2 is down.

Complete the following steps to carry out this scenario.

1. **Start H1->H2->H3**

Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. Once the session is in a Prepared state, both sites 2 and 3 will have a consistent copy of the data

**Note:** Production I/O to H1 can continue during the following steps.

## 2. Unplanned outage of H2

This event causes an unplanned outage of H2, which causes the session to go suspended. It is recoverable if the session was in a Prepared state when the outage occurred. If the session state was Preparing, the session might be recoverable at Site 2 but would not be recoverable at Site 3.

## 3. Start H1->H3

This command starts a Global Mirror session between Site 1 and Site 3 during the outage of H2. Production I/O can continue on H1.

**Note:** If you are running DS8000 devices on H1 with the proper level of microcode, and enough time has passed since the initial start of the session, the creation of the H1 to H3 pairs will NOT require a full copy.

Once the session has reached a Prepared state, a consistent copy of the data will reside on Site 3.

**Note:** Data on Site 2 is no longer consistent at this point.

**Note:** After Site 2 has been successfully brought back up, perform the next step to return to the three-site configuration.

## 4. Start H1->H2->H3

Issue this command when H2 is ready to be brought back into the configuration. This command restores the original configuration. Production I/O can continue on H1. Once the session reaches a Prepared state, a consistent copy will reside on Site 2 and Site 3.

### Unplanned outage of remote (H3):

In this scenario, an unplanned outage occurs on Site 3 (H3).

Complete the follow these steps to carry out this scenario.

#### 1. Start H1->H2->H3

Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. Once the session is in a Prepared state, both sites 2 and 3 will have a consistent copy of the data.

A failure of the remote site (H3) causes the session state to go from Prepared to SuspendedH2H3, with an error message that the connection to the H3 box has been lost.

**Note:** When the remote site (H3) is repaired and ready to return to the configuration, the session will remain SuspendedH2H3 and recoverable until step 2 is completed.

#### 2. Start H1->H2->H3

Issue this command only when the remote site (H3) is repaired and ready to return to the configuration. With this command, your session returns to the original configuration by starting the Global Mirror portion back up.

### Unplanned outage of local (H1) with production move to remote (H3):

In this scenario, there is an unplanned outage for Site 1 (H1). Site 2, however, does not have the resources to run production, so production will be moved to Site 3 during the outage at Site 1. When the outage is complete, production can then be moved back to Site 1 and resume the full three-site configuration.

Complete the following steps to carry out this scenario.

**Note:** To ensure consistency in the event of a rolling disaster where the IBM Tivoli Storage Productivity Center for Replication server might be lost on Site 1 before freezing all the copies in the session, enable the Metro Mirror heartbeat function.

1. **Start H1->H2->H3**

Issue this command to begin Metro Global Mirror on your system when the production I/O is running on H1. When the session is in a Prepared state, both sites 2 and 3 have a consistent copy of the data.

2. **Unplanned outage of H1 and H2**

This event causes an unplanned outage of H1 and H2 which causes the session to go suspended. It is recoverable if the session was in a Prepared state when the outage occurred. If the state was Preparing, the session might be recoverable at Site 2, but will not be recoverable at Site 3.

3. **Recover H3**

This command makes H3 Target Available. The session will be severe/Target Available. When the session state is Target Available at H3, production I/O can be moved to H3.

**Note:** Once Site 1 is brought back up, perform the following step to begin the process of returning production to Site 1.

4. **Start H3->H1->H2**

Issue this command when H1 and H2 are available and ready to be brought back into the configuration. Changes from production at the remote site (H3) flow from the local site (H1) to the intermediate site (H2). The H3-H1 and H1-H2 pairs are Global Copy. The session will in the preparing state.

**Note:** While the session is Preparing, you can decrease the amount of time that production I/O has to be stopped by waiting for the progress on the H3-H1 role pair to be at or near 100%. The next step will drive H3-H1 to sync, which depending on the percentage of the copy that is complete, can take a significant amount of time. When you are ready to switch back to Site 1, stop production I/O on H3 before continuing to next step.

5. **Suspend**

This command suspends the Global Copy relationship between H3>H1 and H1>H2. When this step is complete, the session state should be suspended and recoverable.

6. **Recover H1**

This command makes H1 Target Available. When the session state is Target Available at H1, run production I/O to H1.

7. **Start H1->H2->H3**

This command restores the original three-site configuration. Once the session reaches a Prepared state, a consistent copy of the data will reside on Site 2 and Site 3.

## **Practicing disaster recovery scenarios**

This section provides scenarios for how to practice disaster recovery using a regular Metro Global Mirror session. You can follow these scenarios to practice disaster recovery while continuing to be disaster recovery-capable using only the Metro Mirror portion of the session. The Metro Global Mirror with Practice session is also available, which allows for practicing disaster recovery while maintaining full three-site disaster recovery capability.

**Practice a recovery to H2 with production that is continuing on H1 in a normal configuration:**

In this scenario, your goal is to practice a recovery to H2 with production that is continuing on H1 in a normal configuration.

Follow these steps to carry out this scenario.

**1. Start H1->H2->H3**

Issue this command to begin a Metro Global Mirror session on your system when the production I/O is running on H1. Wait for the session to go to the prepared state.

You can continue to run production I/O to H1 during the following steps.

**2. Suspend**

This command creates the planned outage by suspending the Metro Mirror relationship between H1 and H2. When this step is complete, the session state should be suspended and recoverable.

**Note:** Stop running production I/O to H2 prior to issuing the Suspend command.

**3. Recover H2**

This command makes H2 target-available. When the session is target-available at H2, run practice I/O to H2.

**4. Start H1->H3**

Issue this command to ensure that any I/O that is written to site 1 is copied to site 3.

You can also use the Start H1->H3 command to practice for a disaster at the H2 site while still maintaining disaster recovery capability. When the H1->H3 role pair is in Global Mirror mode, you can recover the session at site 3 if a disaster occurs.

When you complete testing on H2, you can issue the Start H1->H2->H3 command to return to a three-site configuration without completing a full copy.

**5. Start H1->H2->H3**

Issue this command to restart the Metro Global Mirror session in the H1-H2-H3 configuration. Any changes from production at the local site (H1) flow from the remote site (H3) to the intermediate site (H2), which overwrites the test data on H2.

Table 13. Role-pair states at the completion of each step (participating and nonparticipating role pairs)

| Steps | Role-pair states |                   |          |                  |                   |                   |
|-------|------------------|-------------------|----------|------------------|-------------------|-------------------|
|       | H1-H2            | H2-H3             | H2-J3    | H3-J3            | H1-H3             | H1-J3             |
| 1     | Prepared         | Preparing         | Prepared | Target available | Non-participating | Non-participating |
| 2     | Suspended        |                   |          |                  |                   |                   |
| 3     |                  | Non-participating | Prepared |                  |                   |                   |
| 4     |                  |                   |          |                  | Participating     | Prepared          |
| 5     | Prepared         | Preparing         | Prepared |                  | Non-participating | Non-participating |

**Practice recovery to H1 with production that is continuing on H2:**

In this scenario, your goal is to practice a recovery to H1 with production that is continuing on H2.

Follow these steps to carry out this scenario.

1. **Start H1->H2->H3**

Issue this command to begin a Metro Global Mirror session on your system when the production I/O is running on H1. Wait for the session to go to the prepared state.

2. **Suspend**

This command creates the planned outage by suspending the Metro Mirror relationship between H1 and H2. When this step is complete, the session state should be suspended and recoverable.

**Note:** Stop running production I/O to H1 prior to issuing the Suspend command.

3. **Recover H2**

This command makes H2 target-available.

**Note:** Start running production I/O to H2.

4. **Start H2->H1->H3**

This command starts a Metro Global Mirror session with H1 and H2 roles reversed from the original configuration.

You can continue to run production I/O to H2 for the next two steps.

5. **Suspend**

This command creates the planned outage by suspending the Metro Mirror relationship between H1 and H2. When this step is complete, the session should be suspended and recoverable.

6. **Recover H1**

This command makes H1 target-available. When the session state is target-available at H1, run practice I/O to H1.

7. **Start H2->H3**

Issue this command to ensure that any I/O that is written to site 1 is copied to site 3.

You can also use the Start H2->H3 command to practice for a disaster at the H1 site while still maintaining disaster recovery capability. When the H2->H3 role pair is in Global Mirror mode, you can recover the session at site 3 if a disaster occurs.

When you complete testing on H1, you can issue the Start H2->H1->H3 command to return to a three-site configuration without completing a full copy.

8. **Start H2->H1->H3**

Issue this command to restart the Metro Global Mirror session in the H2-H1-H3 configuration. Any changes from production at the local site (H2) flow from the remote site (H3) to the intermediate site (H1), which overwrites the test data on H1.

Table 14. Role-pair states at the completion of each step (participating and nonparticipating role pairs)

| Steps | Role-pair states |                   |                   |                  |                   |                   |
|-------|------------------|-------------------|-------------------|------------------|-------------------|-------------------|
|       | H1-H2            | H2-H3             | H2-J3             | H3-J3            | H1-H3             | H1-J3             |
| 1     | Prepared         | Preparing         | Prepared          | Target available | Non-participating | Non-participating |
| 2     | Suspended        |                   |                   |                  |                   |                   |
| 3     |                  |                   |                   |                  |                   |                   |
| 4     | Prepared         | Non-participating | Non-participating |                  |                   |                   |
| 5     | Suspended        |                   |                   |                  |                   |                   |
| 6     |                  |                   |                   |                  |                   |                   |
| 7     |                  | Preparing         | Prepared          |                  | Non-participating | Non-participating |
| 8     | Prepared         | Non-participating | Non-participating |                  | Preparing         | Prepared          |

### Practice at H3 for Metro Global Mirror:

This topic describes how to practice at the remote site (H3) while running a Metro Global Mirror session.

To practice on the remote site, you first need to suspend the back-end. For example, if running in an H1->H2->H3 configuration, issue a **Suspend H2H3** command to pause the Global Mirror portion of the session while leaving the Metro Mirror portion of the session active. When the session is in a suspended state (that is, H2H3 is suspended while in an H1->H2->H3 configuration), you have the option to issue a **Recover H3** command.

The **Recover H3** command acts differently than other recover commands. In this case, the command acts only to make the remote site (H3) target-available and does not switch the application site to H3. The application site remains the original host site and the session acts as it does in a normal Metro Mirror Failover/Failback session.

While the session is in this mode, it remains capable of disaster recovery at Site 2 and will freeze all pairs to maintain consistency during an error. You can also suspend the session by command and recover to Site 2 if necessary. However, if recovering to Site 2 in this mode, it will not be possible to restart a full three-site solution. To get three sites back up and running, you need to copy back to H1 using the **Start H2->H1** command and then suspend and recover to H1. Once H1 is the application host again, you can issue the **Start H1->H2->H3** command to restart three-site support. If this is done, all data written to H3 during the practice is overwritten with the data that is on H1.

#### *Overview of scenarios for practicing at H3 for Metro Global Mirror:*

This topic provides an overview of the scenarios for practicing at H3 for Metro Global Mirror.

**Note:** For simplicity, the following scenarios are geared toward the configuration where the session is running H1->H2->H3. The same options apply when the configuration is running H2->H1->H3.

There are two methods to set up a practice mode at the remote site, H3.

In the first method, you issue the following commands:

1. **Start H1->H2->H3**
2. **SuspendH2H3**
3. **RecoverH3**

When the session state is Target Available at H3, you can practice your disaster recovery steps on H3 while continuing to run the real production I/O to H1.

In the second method of setting up a practice mode, issue the following commands:

1. **Start H1->H2->H3**
2. **Suspend**
3. **RecoverH3**
4. **Start H1->H2**

When the session state is Target Available at H3, run practice I/O to H3. Continue running production I/O to H1.

The difference between the second method and the first method is that a **Suspend** command is issued instead of a **SuspendH2H3** command. When the session is in suspended mode (as opposed to suspended H2H3 mode), the session does not automatically go into practice mode. The application site switches to H3 and you must issue a **Start H1->H2** to enable the practice mode.

*Practice scenario 1: While practicing, disaster occurs at Site 1, recover to Site 2:*

This topic describes what to do in the event that a disaster occurs on Site 1 while practicing at Site 3.

If a disaster occurs at Site 1, the session immediately suspends itself. You then must complete the following steps:

1. Issue **RecoverH2**.
2. Move the production I/O to Site 2.
3. Issue **Start H2->H1** to make the session capable of disaster recovery again once Site 1 has been brought back up.

*Practice scenario 2: While practicing, planned outage at Site 1, recover to Site 2:*

This topic describes what to do if there is a planned outage on Site 1 while practicing at Site 3.

While practicing, if you want a planned outage with a switch to Site 2, complete the following steps:

1. Issue **Suspend**.
2. Issue **RecoverH2**.
3. Move the production I/O to Site 2.
4. Issue **Start H2->H1** to make the session capable of disaster recovery again once Site 1 has been brought back up.

*Practice scenario 3: while practicing with production at site 2, move production back to site 1:*

This topic describes an H3 practice scenario, with production running at site 2, followed by a move of production to site 1.



While practicing, if you moved the production to site 2, you need to move it back to site 1 before you can start three-site support. To move production back to site 1, do the following steps (assume the session is running H2->H1):

1. Issue **Suspend**.
2. Issue **Recover H1**.
3. Move the production I/O to site 1.
4. Issue **Start H1->H2** to make the session capable of disaster recovery again and move the production I/O to H3.
5. Issue **Start H1->H2->H3** when H3 is ready to receive a copy of the data at site 1. When this command completes, all writes that were done to H3 are overwritten by the data on site 1.

The following table shows the changing states of the role pairs after each step.

*Table 15. Role-pair states at the completion of each step (participating and nonparticipating role pairs)*

| Steps | Role-pair states |                  |          |                  |                  |       |
|-------|------------------|------------------|----------|------------------|------------------|-------|
|       | H1-H2            | H2-H3            | H2-J3    | H3-J3            | H1-H3            | H1-J3 |
| 1     | Suspended        | Nonparticipating |          |                  |                  |       |
| 2     | Target Available |                  |          |                  |                  |       |
| 4     | Prepared         | target-available |          |                  | nonparticipating |       |
| 5     |                  | Preparing        | Prepared | Target Available |                  |       |



---

## Chapter 3. Administering

Administer DB2 to back up the database and to start and stop the control center. Administer IBM Tivoli Storage Productivity Center to authorize users, start and use the console, which is the graphical user interface (GUI), start and stop services, and many other administrative tasks.

---

### Administering DB2

Administer DB2 by backing up the IBM Tivoli Storage Productivity Center database, starting the control center, and starting and stopping DB2.

#### Backing up the Tivoli Storage Productivity Center database

IBM Tivoli Storage Productivity Center uses circular logs for the DB2 database. Therefore, only an offline database backup is supported. When an offline database backup is performed, the database connections are terminated. Therefore, the Tivoli Storage Productivity Center servers need to be stopped and then restarted.

1. Close the Tivoli Storage Productivity Center GUI.
2. Stop the Tivoli Storage Productivity Center servers.
3. Back up your database.
  - a. Open the DB2 Control Center.
  - b. The Control Center View panel is displayed. Select **Basic** and click **OK**.
  - c. In the Control Center navigation tree, click **All Databases**.
  - d. Under the **All Databases** node, you see **TPCDB**. **TPCDB** is the Tivoli Storage Productivity Center database. Highlight **TPCDB** and right-click it. Click **Backup**.
  - e. The Confirm the details of your database panel is displayed. Click **Next**.
  - f. The Specify where to store your backup image panel is displayed. For Media Type, enter **File System**. If you want to add a new folder, click **Add**.
  - g. On the Path Browser - LOCALHOST panel, click the **Create New Folder** icon. Enter the name of your new folder (for example, TPCbackup1). Click **OK**.
  - h. On the Specify where to store your backup image panel, you see your new folder listed (TPCbackup1). Click **Finish**. The database backup runs.
  - i. When the database backup is completed, you see a DB2 Message window indicating that the database backup completed successfully.
  - j. Close the DB2 Message window.
4. Restart the Tivoli Storage Productivity Center servers.

If you want to do an online backup of the database, you need to change the use of circular logs to linear logs. The use of linear logs can fill up the file system quickly if the database is not maintained properly. For information about how to change the use of circular logs to linear logs, see the DB2 Information Center. Click **Designing > Backups**.

## Starting the Control Center

This topic describes how to start the DB2 Control Center.

To start the Control Center, perform the following action:

|            |                                                                                                     |
|------------|-----------------------------------------------------------------------------------------------------|
| On Linux   | From a console window, run the <b>db2cc</b> command.                                                |
| On UNIX    | Open the <b>IBM DB2</b> folder on the desktop and click <b>Control Center</b> .                     |
| On Windows | Click <b>Start &gt; Programs &gt; IBM DB2 &gt; General Administration Tools &gt; Control Center</b> |

## Using the command line on UNIX and Linux

This topic describes how to use a command line to perform actions against a DB2 instance under UNIX or Linux.

If the DB2 control center is unavailable or you do not have access to a graphical user interface, you can use a command line to execute DB2 commands such as starting and stopping an instance. To use a command line to perform actions against an instance of DB2, complete the following steps:

1. Log in with a user ID or name that has **ROOT**, **SYSADM**, **SYSCTRL**, or **SYSMAINT** authority on the instance; or log in as the instance owner.
2. Run the startup script:
  - For Bourne or Korn shell, type: **. HOME/sqllib/db2profile**.
  - For C shell, type: **source HOME/sqllib/db2cshrc**.

where **HOME** is the home directory of the instance you want to use.

3. To start the instance using the command line, type **db2start**.

**Note:** When you run commands to start the database manager instance, the DB2 database manager applies the command to the current instance.

4. To stop the instance using the command line, type **db2stop**.

**Note:** When you run commands to stop the database manager instance, the DB2 database manager applies the command to the current instance.

## Manually starting DB2 on Windows

This topic describes how to manually start DB2 on Windows.

To start DB2 manually, complete the following steps:

1. Start the following Windows services:
  - DB2 - DB2-0
  - DB2DAS - DB2DAS00
  - DB2 JDBC Applet Server
  - DB2 License Server
  - DB2 Security Server
2. To open a DB2 command window, click **Start > Programs > IBM DB2 > Command Line Tools > Command Window**.
3. From the DB2 Command window, run the **db2start** command.

## Manually stopping DB2 on Windows

This topic describes how to manually stop DB2 on Windows.

To stop DB2 manually, complete the following steps:

1. Stop the following Windows services:
  - DB2 Security Server
  - DB2 License Server
  - DB2 JDBC Applet Server
  - DB2DAS - DB2DAS00
  - DB2 - DB2-0

**Note:** When you stop the **DB2 Security Server** service, you are prompted to stop the Warehouse logger and Warehouse Serve; click **Yes**.

2. To open a DB2 command window, click **Start > Programs > IBM DB2 > Command Line Tools > Command Window**.
3. From the DB2 Command window, run the **db2stop** command.

To restart DB2, enter the **db2start** command from the DB2 Command window. Before you can issue the command, you must first start these services from the Windows Services panel:

```
DB2 - DB2-0
DB2DAS - DB2DAS00
DB2 JDBC Applet Server
DB2 License Server
DB2 Security Server
```

## Monitoring DB2

The minimum user authority level needed for monitoring DB2 instances is a user with DB2 system maintenance authority (SYSMAINT).

To check and set SYSMAINT authority, follow these steps:

1. Run this command in the DB2 command prompt window to check to see if there is an operating system user group defined to have SYSMAINT authority:  
`db2 get dbm cfg`

In the output file, look for this information:

```
SYSADM group name (SYSADM_GROUP) =
SYSCTRL group name (SYSCTRL_GROUP) =
SYSMAINT group name (SYSMAINT_GROUP) =
SYSMON group name (SYSMON_GROUP) =
```

If the setup for the operating system group has not been done, you do not see a value set.

If the setup has been done, this example shows what you can expect to see:

```
SYSADM group name (SYSADM_GROUP) =
SYSCTRL group name (SYSCTRL_GROUP) =
SYSMAINT group name (SYSMAINT_GROUP) = ADMINISTRATORS
SYSMON group name (SYSMON_GROUP) =
```

In this example, the "ADMINISTRATORS" group has SYSMAINT\_GROUP authority.

2. If the setup has been done, add the user you want to use to the ADMINISTRATORS group using the operating system utilities or use a user that already belongs to the ADMINISTRATORS group.  
If you want to give a user group "SYSMAINT\_GROUP" authority, follow these steps:
  - a. If a user (for example userA) belongs to an operating system group called db2monitor, here is an example of setting the db2monitor group with SYSMAINT authority. From the DB2 command prompt window, run the following command:

```
db2 update dbm cfg using SYSMAINT_GROUP db2monitor
```
  - b. After issuing the **db2 update** command, restart DB2 by running the following command from the DB2 command prompt window or restarting the system:

```
db2 force application all
```

This command might need to be issued a few times to stop all the database connections.
  - c. Run the following commands from the DB2 command prompt window:

```
db2stop
db2start
```
  - d. UserA can now monitor the DB2 database.

---

## Adding a Tivoli Storage Productivity Center user group name to Tivoli Storage Productivity Center for Replication

To access Tivoli Storage Productivity Center for Replication using a user ID from the Tivoli Storage Productivity Center user group name, you must add the user group name to the Tivoli Storage Productivity Center for Replication administrative role. If you installed Tivoli Storage Productivity Center with LDAP authentication, you do not have to perform these steps.

Follow these steps:

1. Log in to Tivoli Storage Productivity Center as the Superuser.
2. Expand **Administrative Services > Configuration**. Left-click **Role-to-Group Mappings**. Locate the user group name that has the user ID you want for logging in to Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication.
3. Add the user group name to Tivoli Storage Productivity Center for Replication Administrator role by using the Add Access wizard. Open Tivoli Storage Productivity Center for Replication using a Tivoli Storage Productivity Center for Replication Administrator user ID.
4. Click **Add Access** on the Administration page.
5. Add the user group name to Tivoli Storage Productivity Center for Replication and complete the wizard.

---

## Accessing the Tivoli Storage Productivity Center for Replication GUI

After IBM Tivoli Storage Productivity Center for Replication has been installed, you can access the IBM Tivoli Storage Productivity Center for Replication GUI.

If you are running IBM Tivoli Storage Productivity Center for Replication on AIX, Linux, or Windows, be sure you use IBM Tivoli Storage Productivity Center Standard Edition to install IBM Tivoli Integrated Portal and IBM Tivoli Storage Productivity Center for Replication.

Go to the following URL to access the IBM Tivoli Storage Productivity Center for Replication GUI (the Web address is case-sensitive):

`https://<HostName>:<port>/CSM`

**Note:**

1. The default port is 3443 for open systems and 9443 for z/OS. For information about z/OS, see *IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide*.
2. If you modified the default port settings during the installation, replace 3443 with the port number you selected for the https port.
3. If you are using IBM System Services Runtime Environment for z/OS, the default port is 32200.
4. If you are using IBM Tivoli Storage Productivity Center for Replication version 3.3 or earlier, the default port for open system is 9443.

Use the same user ID and password that you entered during the installation of IBM Tivoli Integrated Portal and IBM Tivoli Storage Productivity Center for Replication. .

---

## Backing up and recovering all the Tivoli Storage Productivity Center for Replication

There are three backup methods for IBM Tivoli Storage Productivity Center for Replication. This section provides information about these methods.

The backup methods are:

- The first backup option is to use IBM Tivoli Storage Productivity Center for Replication's import and export function to save the configuration of each IBM Tivoli Storage Productivity Center for Replication session. You would export a IBM Tivoli Storage Productivity Center for Replication session to a CSV file. The CSV file can then be used to recover a IBM Tivoli Storage Productivity Center for Replication session.
- The second backup option is to use IBM Tivoli Storage Productivity Center for Replication's high availability function. This option allows IBM Tivoli Storage Productivity Center for Replication to have a real time backup server that can take over in the event of a disaster.
- The third backup option is to use the **mkbackup** command to create a backup copy of the embedded database repository. This option does not work for DB2 databases. Any user on DB2 (z/OS or IBM Tivoli Storage Productivity Center for Replication before version 4.2) will not be able to use the **mkbackup** command.

For information about the high availability function and configuring sessions and copy sets, go to the Information Center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp>.

---

## Administering Tivoli Storage Productivity Center

This topic provides information about how to administer IBM Tivoli Storage Productivity Center.

### Checking the Tivoli Storage Productivity Center for Replication version

This section describes how to check the version of Tivoli Storage Productivity Center for Replication that is installed on your system.

To check the version of Tivoli Storage Productivity Center you have installed, go to the following directory:

Windows: C:\Program Files\IBM\TPC  
UNIX or Linux: <usr or opt>/IBM/TPC

This directory contains the `version.txt` file which displays the version of Tivoli Storage Productivity Center you have installed.

This is an example of the information in `version.txt`:

```
TPC V3.3.2 Build TPC_3.3.2.56 20080403
```

### Starting and stopping the Tivoli Storage Productivity Center GUI

This topic describes how to start and stop the Tivoli Storage Productivity Center graphical user interface (GUI).

#### Starting the Tivoli Storage Productivity Center GUI

Learn how to start the Tivoli Storage Productivity Center graphical user interface (GUI) from one of the following locations: the Microsoft Windows Start menu, IBM Tivoli Integrated Portal, the Productivity Center icon on your desktop (Windows), or the command line (AIX, UNIX, or Linux).

#### Starting the Tivoli Storage Productivity Center GUI from Tivoli Integrated Portal:

You can start the Tivoli Storage Productivity Center graphical user interface (GUI) from Tivoli Integrated Portal. With the single sign-on feature, you can log on to the Tivoli Storage Productivity Center GUI without having to explicitly enter a user name and password.

To start the Tivoli Integrated Portal and Tivoli Storage Productivity Center, complete the following steps:

1. Ensure that you are using these browsers:

**AIX** Firefox 2.0 or 3.0

**Linux and UNIX**  
Firefox 2.0 or 3.0

**Windows**  
Internet Explorer 7 or 8

Firefox 2.0 or 3.0

2. Ensure that a fully qualified domain name is defined for the Tivoli Storage Productivity Center server. To verify the Tivoli Storage Productivity Center server name and modify the name if required, go to the information center.



Search for **Data server and Device server are inaccessible from Tivoli Integrated Portal**. Contact your Tivoli Storage Productivity Center administrator before you modify the file.

3. If you are using Internet Explorer 8 the first time, click **Start Tivoli Storage Productivity Center**.
4. Click the message and select **Download file**. You can then start Java Web Start.
5. To start Tivoli Storage Productivity Center from Tivoli Integrated Portal, complete the following steps:
  - a. Start an Internet Explorer or Firefox Web browser and type the following information in the address bar:  
`http://hostname:port`

where *hostname* defines the server that is running Tivoli Integrated Portal such as the server name or IP address and *port* defines the port number for Tivoli Integrated Portal. If the default port was accepted during the installation of Tivoli Integrated Portal, the port number is 16310. Contact your Tivoli Storage Productivity Center administrator if you need to verify the host name and port number

- b. On the Tivoli Integrated Portal logon page, log on using the appropriate user ID and password. Your user ID must have administrator permissions.
- c. In the Tivoli Integrated Portal navigation tree, click **Tivoli Storage Productivity Center**
- d. On the Tivoli Storage Productivity Center portlets page, click **Start Storage Productivity Center**, and one of the following occurs:
  - If single sign-on is successful, Tivoli Storage Productivity Center starts without displaying the logon window.
  - If single sign-on is not successful, an error message and the Tivoli Storage Productivity Center logon window are displayed.
  - If the Tivoli Storage Productivity Center Data server or Device server is not accessible, an error message is displayed. The status of both the Data server and the Device server also display.
  - If the status of the Device server and Data server is inaccessible and you started Tivoli Integrated Portal from a remote computer, it is possible that a fully qualified domain name was not defined for the Tivoli Storage Productivity Center server during installation. To check the Tivoli Storage Productivity Center server name and modify the name if required, see the Tivoli Storage Productivity Center troubleshooting documentation.
  - If you are using a Lightweight Directory Access Protocol (LDAP) compliant directory for Tivoli Storage Productivity Center user authentication and the directory is not available, an error message is displayed.

Once you have logged on to Tivoli Integrated Portal, a Lightweight Third-Party Authentication (LTPA) token is created. This token is passed to other applications that you start from Tivoli Integrated Portal for single sign-on authentication purposes.

During the period between when you log on to Tivoli Integrated Portal and when you start another application such as Tivoli Storage Productivity Center from Tivoli Integrated Portal, the following conditions might occur:

- The user password that was used to log on to Tivoli Integrated Portal is changed in the user repository.

- The user ID that was used to access Tivoli Integrated Portal is changed in the repository or removed from the user repository.
- The user repository is not accessible.

Under the first condition, the original user credentials that were used to access Tivoli Integrated Portal are used to access other applications until the timeout period for the LTPA token that is used for single sign-on expires. When the LTPA token expires, you are prompted to re-enter your user ID and password when you attempt to start another application using single sign-on.

Under the second and third conditions, the single sign-on feature does not work. You are always prompted to re-enter your user ID and password when you attempt to start another application.

### **Starting the Tivoli Storage Productivity Center GUI as a stand-alone application:**

This section describes how to start the IBM Tivoli Storage Productivity Center from the Microsoft Windows Start menu, or the Productivity Center icon on your desktop (Windows), or from the command line (AIX, UNIX, or Linux).

To start the Tivoli Storage Productivity Center graphical user interface (GUI) on Windows, click **Start > Programs > IBM Tivoli Storage Productivity Center > Productivity Center**. You can also double-click the IBM Tivoli Storage Productivity Center icon if it is installed on your desktop.

To start the Tivoli Storage Productivity Center GUI on UNIX or Linux, type the following path and command at the command line:

```
/opt/IBM/TPC/gui/TPCD.sh
```

To start the Tivoli Storage Productivity Center GUI on AIX, type the following path and command at the command line:

```
/opt/IBM/TPC/gui/TPCD.sh
```

### **Logging on to Tivoli Storage Productivity Center:**

Use the logon window to specify the user ID and password that you want to use to log on to a IBM Tivoli Storage Productivity Center server.

Define logon information for Tivoli Storage Productivity Center in the following fields:

#### **User ID**

Enter the user ID that you want to use to log on to Tivoli Storage Productivity Center. The roles that are assigned to that user ID determine what nodes in the navigation tree you can see and act upon.

#### **Password**

Enter the password that is associated with the **User ID**.

**Server** Enter the IP address or Domain Name System (DNS) name of the computer on which the Tivoli Storage Productivity Center server is installed.

When you enter an IP address, you can use either an IPv4 or IPv6 address format depending on the protocol enabled on the computer. For example, enter an IPv6 address for an IPv6-only computer. Enter an IPv4 or IPv6 address for a computer with dual stacks (both IPv4 and IPv6 enabled).

- To enter an IPv4 address, use the dotted decimal format: *nnn.nnn.nnn.nnn*. For example: 127.0.0.1.
- When entering an IPv6 address you must include brackets [] to separate those addresses from their port numbers when you enter a value in this field. For example: [2001:0DB8::1234:0000:0000:5678:ABCD]:9550. Use one of the following methods to enter an IPv6 address:
  - The preferred IPv6 address representation is written as eight groups of four hexadecimal digits: *xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx*, where each *x* is a hexadecimal digit representing 4 bits. For example: 2001:0DB8:0000:1234:0000:0000:5678:ABCD
  - You can specify an IPv6 address using a shortened format that omits leading zeros: (2001:DB8:0:1234:0:0:5678:ABCD)
  - You can use double colons in place of a series of zeros: (2001:DB8:0000:1234::5678:ABCD)

### Setting up a user role and collecting system statistics

If you are logging on to the Tivoli Storage Productivity Center graphical user interface (GUI) for the first time after installing the product, note the following information:

- If you did not set up a Tivoli Storage Productivity Center user role, you might not be able to log on. To set up a Tivoli Storage Productivity Center user role, log on as the superuser and create the roles as described in the IBM Tivoli Storage Productivity Center Information Center.
- After you log on, you will see a panel that prompts you to collect system statistics. Click **Yes**.

### Single sign-on

Tivoli Storage Productivity Center can use the single sign-on feature, which enables you to start Tivoli Storage Productivity Center from either the IBM Tivoli Integrated Portal instance that is installed with Tivoli Storage Productivity Center or from an existing Tivoli Integrated Portal V1.1.1.x instance. If you have a single sign-on environment configured and Tivoli Integrated Portal user authentication is successful, the Tivoli Storage Productivity Center logon window does not open when you start the Tivoli Storage Productivity Center GUI from Tivoli Integrated Portal.

### User credential considerations for single sign-on

When you log on to either Tivoli Integrated Portal or Tivoli Storage Productivity Center, a Lightweight Third-Party Authentication (LTPA) token is created and used for single sign-on authentication when you start other applications from within the Tivoli Storage Productivity Center GUI.

During the period between when you log on to Tivoli Storage Productivity Center and when you start other applications, such as element managers, from the Tivoli Storage Productivity Center GUI, the following conditions might occur:

|                                                                                                                                                |                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The user password that was used to log on to Tivoli Integrated Portal or Tivoli Storage Productivity Center is changed in the user repository. | Under this condition, the original user credentials that were used to log into Tivoli Integrated Portal or Tivoli Storage Productivity Center are used to access other applications until the timeout period for the LTPA token that is used for single sign-on expires. When the LTPA token expires, you are prompted to reenter your user ID and password when you attempt to start another application using single sign-on. |
| The user ID that was used to access Tivoli Integrated Portal is changed in the repository or removed from the user repository.                 | Under the second and third conditions, the single sign-on feature does not work. You are always prompted to re-enter your user ID and password when you attempt to start another application through single sign-on.                                                                                                                                                                                                            |
| The user repository is not accessible.                                                                                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## Stopping the Tivoli Storage Productivity Center GUI

This section describes how to stop the IBM Tivoli Storage Productivity Center graphical user interface (GUI).

To stop the Tivoli Storage Productivity Center user interface, click **File > Exit**.

## Viewing and managing services

This topic describes viewing and managing services for Data server and Device server.

### Data Server services

Data Server directs all monitoring operations. Data Server services include the various services that run in the IBM Tivoli Storage Productivity Center for Data Server component.

The Data Server is the control point for scheduling functions, configuration, event management, reporting, and GUI support. The Data Server directs all of the activities related to the monitoring performed on computers with agents installed and the reporting performed on the client GUI computers.

When you expand the Data Server node, five nodes appear which represent the running services within the server. Select any of the following nodes to display information on the server identification, run time statistics, requests processed, database information, and agent statistics:

**Server** The component that provides the functionality of the Data Server. The Server, sometimes called Data Server, provides monitoring functions for services, data, and facilities to the network.

#### Scheduler

The Scheduler is the process that receives requests to schedule operations, such as the running of scheduled pings, probes, and scans. It is a multithread background server process designed to handle the scheduling and launching of jobs, based on a simple timing scheme.

**GUI** The main graphical user interface that enables you to do monitoring operations, such as generate reports, check the status of agents, create and run jobs, setup monitoring for the Data Server, and so forth.

**Device**

Converts general commands for peripheral devices and passes them between the Server and the device. Defines I/O devices to the operating system and then controls the operation of these devices.

**Agent** The process on the server that sends and receives requests from agents for information gathering or processing. Agents can perform specific operations, such as scans or probes, on a regular schedule.

With Data Server services you can perform the following tasks:

- View a Data Server service log
- Configure tracing on a service
- Shut down a service
- Restart a service that is inactive

Additionally, through the Server service you can do the following:

- Configure audit logging
- Broadcast name and port information

To view the status of the server and services, expand **Administrative Services > Services > Data Server**. A green square icon next to the name of a service (for example, **Scheduler**) means that the service is active and running normally. If the service is not currently active, a red circle is displayed next to the name.

**Refresh the Data Server list**

To refresh the Data Server list, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services > Services**.
2. Right-click **Data Server**, and click **Refresh**.

**Viewing Data Server information with the Server service:**

The Server service directs all Data Server information. You can view information about the Server service, including the name and port, disk space, connections, requests, and agent information.

To view Data Server information, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services > Services > Data Server**.
2. Click **Server**. In the right pane the Server service page opens displaying the following information:

**Server** Displays the name of the server.

**Port** Displays the port number on which the server is listening to agents.

**Start Time**

Displays the date and time when the server was started.

**Elapsed Time**

Displays the elapsed time when the server was last updated.

**VM Size**

Displays the disk space allocated to the virtual memory size of the server.

**Network Connections**

Displays the number of networks to which the server is connected.

**Maximum Connections**

Displays the maximum number of networks to which the server can connect.

**Pending Requests**

Displays the number of outstanding requests waiting for server response.

**Processed Requests**

Displays the number of requests processed by the server.

**Database**

The following fields display information about the Data Manager's repository database.

**JDBC URL**

Displays the location of the JDBC file used to connect to the database.

**JDBC Driver**

Displays the name of the JDBC driver used to connect to the database.

**DB Connections**

Displays the number of connections to the database.

**DB Pool Count**

Displays the database pool count.

**Agents**

The following fields provide information about the status of Data Manager agents.

**Total** The number of agents that are installed.

**Online**

The number of agents that are currently online (communicating with the server).

**Offline**

The number of agents that are currently offline (not communicating with the server, but were shut down normally).

**Unreachable**

The number of agents that are 1) not communicating with the server, and 2) provided no information to the server that they were shutting down.

**Upgrading**

The number of agents that are currently in the process of upgrading to the same version as the server.

**Upgrade Required**

The number of agents that must be upgraded to the same version as the server.

**Viewing information about Data Server services:**

Select any of the Data Server services to display information on that services' identification, run time statistics, requests processed, database information, and agent statistics.

To display information about a service, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services > Services > Data Server**.
2. Click one of the following services:

**Scheduler**

Controls the running of scheduled pings, probes, and scans.

**GUI** Enables you to generate report requests and to create and run jobs.

**Device**

Converts commands for peripheral devices.

**Agent** Sends and receives requests from agents.

3. In the right pane, the following information about the selected Data server service is displayed:

**Service**

Whether the service is active or down.

**Pending Requests**

Number of requests waiting to be processed by the scheduler for the current session.

**Processed Requests**

Number of requests processed by the scheduler for the current session.

**Idle Threads**

Number of threads waiting for a process request.

**Viewing a Data Server service log:**

Use the content of a service log file to troubleshoot any errors that might occur during startup, processing, or shutdown of the Data Server service.

To display a service log, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services > Services > Data Server**.
2. Right-click the name of the service and click **View Log**. The service log will open in the right pane.

Service alerts, if any, are displayed in red.

Logs are searchable by search string. Enter the string at the top of the page and click the up or down arrow to search.

**Configuring tracing for Data Server services:**

Configure tracing for the Scheduler, GUI, Device, and Agent services by setting the debug level. Tracing changes can only be made by users with administrator authority.

To configure tracing for the Scheduler, GUI, Device, and Agent services, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services > Services > Data Server**.
2. Right-click the name of the service and click **Configure Tracing**.

3. In the Server Tracing Configuration window, check the **Enable Trace** box to enable tracing for the service. If this box is unchecked, tracing will not be performed.
4. Once tracing has been enabled, select the level of debugging you want to perform. You can select one of the following:
  - **DEBUG\_MIN** - Provides minimal information on errors. This is the default value.
  - **DEBUG\_MID** - Provides more information on errors.
  - **DEBUG\_MAX** - Lists all debugging trace information.
5. Click **OK**.

Tracing can be turned off by de-selecting the **Enable Trace** option in the Server Tracing Configuration window.

### **Configuring audit logging for the Server:**

Use audit logging to record a history of user actions. This function is available for the Server only.

To configure audit logging, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services > Services > Data Server**.
2. Right-click the **Server** node and click **Configure Audit Logging**. The Server Audit Logging Configuration window opens.
3. In the Server Audit Logging Configuration window, check the **Enable Trace** check box to enable tracing for the server component. If this box is unchecked, tracing will not be performed.
4. Specify the audit logging options:

#### **Maximum Number of Files**

Select the maximum number of trace files that will be created. When the maximum number of files has been reached, tracing will rollover and start writing to the first file. (default is 5)

The value in this field also applies to the other services for which tracing is enabled.

#### **Maximum File Size (kilobytes)**

Select the maximum size of each log file. When the current trace file has reached its maximum capacity, tracing will continue to the next available trace file as defined in the Maximum Number of Files field. (default is 20 MB).

The value in this field also applies to the other services for which tracing is enabled.

5. Click **OK**.

### **Broadcasting the Data Server location:**

Use broadcasting to inform agents of the Data Server location. Broadcasting is available with the Server service only.

Broadcasting informs agents of the server location

To broadcast the Data Server location, complete the following steps:



1. In the Navigation Tree pane, expand **Administrative Services > Services > Data Server**.
2. Right-click the **Server** node and click **Broadcast**.

### Shutting down a Data Server service:

Shut down the active process or processes associated with a service (for example, processes on the server that receive requests from the scheduler). If you shut down a service, you must restart it before the server will resume receiving requests from it.

To shut down a running service, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services > Services > Data Server**.
2. Right-click the name of the service and click **Shutdown**.
3. Select one of the following shutdown methods:

#### Normal

Let all running and scheduled processes complete and then shut down cleanly. During this shutdown process, new jobs are allowed to run if they were scheduled during the time that the shutdown is taking place. Once all jobs have completed and no more jobs are scheduled to start during the shutdown time, the service will be shut down.

#### Immediate

Let all running processes complete but not allow scheduled jobs to start. During this shutdown process, new jobs are *not* allowed to start if they were scheduled during the time that the shutdown is taking place. Once all running jobs have completed, the service will be shut down.

**Abort** Stops whatever processes are currently running and then shuts down the service.

### Starting a Data Server service:

Reactivate a service after it has been stopped. The start action is available only for services that are not running.

To start a service, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services > Services > Data Server**.
2. Right-click the name of the service and click **Start**.

### Configuring tracing

Configure tracing for the Data server, agent, scheduler, Device server, and GUI by setting the debug level and log file characteristics. The level of information saved in a log file is determined by the tracing level. Tracing settings can only be changed by users with administrator authority.

To configure tracing for IBM Tivoli Storage Productivity Center, complete the following steps:

1. In the navigation tree pane, expand **Administrative Services > Services > Data Server**.
2. Right-click one of the following components:
  - Server

- Agent
- Scheduler
- Device
- GUI

Click **Configure Tracing**.

3. In the Server Tracing Configuration window, check the **Enable Trace** check box. Specify additional tracing options:

**Level** Select the level of debugging you want to perform. Select one of the following debugging level:

**DEBUG\_MIN**

Provides minimal information about errors. This option is the default value.

**DEBUG\_MID**

Provides more information about errors.

**DEBUG\_MAX**

Lists all debugging trace information.

**Maximum number of files**

Select the maximum number of trace files that is created. When the maximum number of files has been reached, tracing rolls over and start writing to the first file. (default is five)

The value in this field also applies to the other services for the Data server for which tracing is enabled.

**Maximum File Size (kilobytes)**

Select the maximum size of each log file. When the current trace file has reached its maximum capacity, tracing continues to the next available trace file as defined in the Maximum Number of Files field. (default is 20 MB).

4. Click **OK**.

Tracing can be turned off by clearing the **Enable Trace** option in the Server Tracing Configuration window.

**Device server services**

A **Device server** node is present in the Navigation Tree pane only if a device server is deployed in IBM Tivoli Storage Productivity Center your environment.

In the Navigation Tree pane, expand **Administrative Services > Services > Device server** to display a list of the services associated with the Device server. The following services are displayed:

**Device server**

Detects all the operations of the Device server.

**SAN Planner**

Monitors the SAN planning tools.

**Event** Service that monitors external events.

**Log** Service for configuring logging and tracing.

**Configuration Analysis and History**

Monitors the SAN configuration tools.

- Control**  
Service that provides control functions.
- Fabric Manager**  
Service that provides fabric management functions.
- Disk Manager**  
Service that provides storage subsystem management functions.
- Inband Agent**  
Service that maintains the list of in-band Fabric agents.
- Disk Monitor**  
Service that monitors the storage subsystems.
- Out of Band Agent**  
Service that maintains the list of out-of-band Fabric agents.
- Element Manager**  
Service that provides ability to manage element managers.
- Performance**  
Service that monitors device performance.
- Configuration**  
Service for configuring Device server properties.
- Discovery**  
Service that provides discovery and probe functions.
- Tape Manager**  
Service that provides tape management functions.
- Message**  
Service that provides internal messaging functions.
- Data Access functions**  
Service providing data access and control functions.
- Data Server Monitor**  
Service that provides Data server monitoring functions.
- Service Manager**  
Service that manages the other services.
- Server Virtualization Management Service**  
Service providing server virtualization management functions.
- Device server Monitor**  
Service that provides Device server monitoring functions.

A green square icon next to a service name means that the service is active and running normally. If the service is not currently active, a red circle icon displays next to the name.

#### **Changing the Device server password:**

This section provides information about changing the Device server password.

To change the password, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services > Services > Device server**.
2. Right-click the Device server node and select **Change Password**.

3. In the Change Password window, type the name of the new password twice and click **OK**.

### **Displaying information about a Device server service**

You can view service information including start time, version, and status.

To display information about a service, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services > Services > Device server**.
2. Click the name of the service.

### **Viewing a Device server log**

You can use the content of the log file to troubleshoot any errors that might occur during startup, processing or shut down. Service alerts, if any, are displayed in red.

To display the log of a service, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services > Services > Device server**.
2. Right-click the name of the service and select **View Log**.

## **Starting and stopping the Tivoli Storage Productivity Center services**

This section provides information about how to start and stop the IBM Tivoli Storage Productivity Center services.

**Note:** The IBM Tivoli Integrated Portal server starts automatically after you reboot the system or after a temporary power loss on Windows. The IBM Tivoli Integrated Portal server does not start automatically on AIX or Linux; you must manually start the server.

### **Stopping the IBM Tivoli Storage Productivity Center services**

This topic describes how to stop the IBM Tivoli Storage Productivity Center services.

#### **Stopping the services on Windows**

To stop the services on Windows, click **Start > Control Panel > Administrative Tools > Services**. Stop the following services:

**IBM WebSphere Application Server V6.1 - DeviceServer**

**IBM Tivoli Storage Productivity Center - Data Server**

**IBM Tivoli Common Agent *directory***

where *directory* is where the common agent is installed. The default directory is c:\Program Files\IBM\TPC\ca. Tivoli Common Agent Services stops the Common agent and any Data agent or Fabric agent running under it.

**IBM Storage Resource Agent - *directory***

where *directory* is where the Storage Resource agent is installed. The default directory is c:\Program Files\IBM\TPC\agent.

**IBM WebSphere Application Server V6.1 - CSM**

**Tivoli Integrated Portal - TIPProfile\_Port\_XXXXX**

where XXXXX indicates the port specified during installation.

## IBM ADE Service

IBM ADE Service is the Tivoli Integrated Portal registry.

## Stopping the services on Linux

To stop the services on Linux, go to a command prompt window and enter the following commands:

### Data server

```
/TPC_install_directory/data/server/tpcdsrv1 stop
```

The default directory is */opt/IBM/TPC*.

### Device server

```
/TPC_install_directory/device/bin/linux/stopTPCF.sh
```

The default directory is */opt/IBM/TPC*.

### Common agent

```
/common_agent_install_directory/ca/endpoint.sh stop
```

The default directory is */opt/IBM/TPC*.

### Storage Resource agent

```
/SRA_install_directory/agent/bin/agent.sh stop
```

The default directory is */opt/IBM/TPC*.

## IBM WebSphere Application Server V6.1 - CSM

```
/usr_or_opt/IBM/replication/eWAS/profiles/CSM/bin/stopServer.sh server1
```

## Tivoli Integrated Portal

```
/install_directory/tip/profiles/TIPProfile/bin/stopServer.sh server1
```

## IBM ADE Service

Follow these steps:

1. Source the environment:  

```
. /var/ibm/common/acsi/setenv.sh
```
2. Run this command:  

```
/usr/ibm/common/acsi/bin/acsisrv.sh -stop
```
3. Wait for the operation to complete.

## Stopping the services on AIX

To stop the services on AIX, go to the command prompt window and enter the following commands:

### Data server

```
stopsrc -s TSRMSrv1
```

### Device server

```
/TPC_install_directory/device/bin/aix/stopTPCF.sh
```

The default directory is */opt/IBM/TPC*.

### Common agent

```
/common_agent_install_directory/ca/endpoint.sh stop
```

The default directory is */opt/IBM/TPC*.

### Storage Resource agent

```
/SRA_install_directory/agent/bin/agent.sh stop
```

The default directory is /opt/IBM/TPC.

### **IBM WebSphere Application Server V6.1 - CSM**

```
/usr_or_opt/IBM/replication/eWAS/profiles/CSM/bin/stopServer.sh server1
```

### **Tivoli Integrated Portal**

```
/install_directory/tip/profiles/TIPProfile/bin/stopServer.sh server1
```

### **IBM ADE Service**

Follow these steps:

1. Source the environment:  
    . /var/ibm/common/acsi/setenv.sh
2. Run this command:  
    /usr/ibm/common/acsi/bin/acsisrv.sh -stop
3. Wait for the operation to complete.

## **Starting the Tivoli Storage Productivity Center services**

This topic describes how to start the Tivoli Storage Productivity Center services.

### **Starting the services on Windows**

To start the services on Windows, click **Start > Control Panel > Administrative Tools > Services**. Start the following services:

#### **IBM WebSphere Application Server V6.1 - DeviceServer**

#### **IBM Tivoli Storage Productivity Center - Data Server**

#### **IBM Tivoli Common Agent *directory***

where *directory* is where the common agent is installed. The default directory is c:\Program Files\IBM\TPC\ca. Tivoli Common Agent Services starts the Common agent and any Data agent or Fabric agent running under it.

#### **IBM Storage Resource Agent - *directory***

where *directory* is where the Storage Resource agent is installed. The default directory is c:\Program Files\IBM\TPC\agent.

#### **IBM WebSphere Application Server V6.1 - CSM**

#### **Tivoli Integrated Portal - TIPProfile\_Port\_XXXXX**

where XXXXX indicates the port specified during installation.

#### **IBM ADE Service**

IBM ADE Service is the Tivoli Integrated Portal registry.

### **Starting the services on Linux**

To start the services on Linux, go to a command prompt window and enter the following commands:

#### **Data server**

```
/TPC_install_directory/data/server/tpcdsrv1 start
```

The default directory is /opt/IBM/TPC.

#### **Device server**

```
/TPC_install_directory/device/bin/linux/startTPCF.sh
```

The default directory is /opt/IBM/TPC.

#### **Common agent**

```
/common_agent_install_directory/ca/endpoint.sh start
```

The default directory is /opt/IBM/TPC.

#### **Storage Resource agent**

```
/SRA_install_directory/agent/bin/agent.sh start
```

The default directory is /opt/IBM/TPC.

#### **IBM WebSphere Application Server V6.1 - CSM**

```
/usr_or_opt/IBM/replication/eWAS/profiles/CSM/bin/startServer.sh server1
```

#### **Tivoli Integrated Portal**

```
/install_directory/tip/profiles/TIPProfile/bin/startServer.sh server1
```

#### **IBM ADE Service**

Follow these steps:

1. Source the environment:  

```
. /var/ibm/common/acsi/setenv.sh
```
2. Run this command:  

```
/usr/ibm/common/acsi/bin/acsisrv.sh -start
```
3. Wait for the operation to complete.

### **Starting the services on AIX**

To start the services on AIX, go to the command prompt window and enter the following commands:

#### **Data server**

```
startsrc -s TSRMsrv1
```

#### **Device server**

```
/TPC_install_directory/device/bin/aix/startTPCF.sh
```

The default directory is /opt/IBM/TPC.

#### **Common agent**

```
/common_agent_install_directory/ca/endpoint.sh start
```

The default directory is /opt/IBM/TPC.

#### **Storage Resource agent**

```
/SRA_install_directory/agent/bin/agent.sh start
```

The default directory is /opt/IBM/TPC.

#### **IBM WebSphere Application Server V6.1 - CSM**

```
/usr_or_opt/IBM/replication/eWAS/profiles/CSM/bin/startServer.sh server1
```

#### **Tivoli Integrated Portal**

```
/install_directory/tip/profiles/TIPProfile/bin/startServer.sh server1
```

#### **IBM ADE Service**

Follow these steps:

1. Source the environment:  

```
. /var/ibm/common/acsi/setenv.sh
```
2. Run this command:  

```
/usr/ibm/common/acsi/bin/acsisrv.sh -start
```
3. Wait for the operation to complete.

## Starting and stopping the Tivoli Integrated Portal server

The IBM Tivoli Integrated Portal server starts automatically after you reboot the system or after a temporary power loss on Windows. The server does not start automatically on AIX or Linux; you must manually start Tivoli Integrated Portal. You can manually stop and start the server before beginning certain configuration tasks or as needed.

Perform the following steps at the command prompt:

1. Change to the following directory:

For Windows:

```
cd TIP_installation_directory\tip\profiles\TIPProfile\bin\
```

For AIX or Linux:

```
cd TIP_installation_directory/tip/profiles/TIPProfile/bin
```

2. To stop the server, enter the following command:

For Windows:

```
stopServer server1
```

For AIX or Linux:

```
./stopServer.sh server1
```

Wait for the server to complete the operation.

3. To start the server, enter the following command:

For Windows:

```
startServer server1
```

For AIX or Linux:

```
./startServer.sh server1
```

Wait for the server to complete the operation.

## Checking Tivoli Storage Productivity Center status

This section provides information on checking the status of Tivoli Storage Productivity Center.

### Checking the status of the Data Server

This topic describes how to check the status of the Data server.

To check the status of the Data Server, complete the following steps:

1. Open the IBM Tivoli Storage Productivity Center graphical user interface (GUI).
2. In the left pane, click **Administrative Services > Services > Data Server**. A status icon is displayed to the left of each service.

### Checking the status of the Device server

This topic describes how to check the status of the Device server.

To check the status of the Device server, complete the following steps:

1. Open the IBM Tivoli Storage Productivity Center graphical user interface (GUI).
2. In the left pane, click **Administrative Services > Services > Device server**. A status icon is displayed to the left of each service.

### Checking the status of agents

This topic describes how to check the status of the agents.



In the navigation tree, expand **Administrative Services > Data Sources**. Click one of the following nodes:

- **CIMOM Agents**
- **Data/Storage Resource Agents**
- **Inband Fabric Agents**
- **Out of Band Fabric Agents**

In the content pane, a list of agents is displayed. If the agents are up and running, there will be a green health status icon in the State column. For CIMOM Agents, there will be a green health status icon in the Connection Status column.

If the status icon is red, you can view the log by selecting an agent and click **View Log** for Data, Native, and Inband Fabric. This action displays the log in the content pane. If the reason that an agent is down is a connectivity problem, you might not be able to get the log files. A given discovered CIMOM appears as a red (critical) icon status until you enter valid credentials.

## Increasing memory allocation

If you run out of memory for the Data server, Device server, or agent, you might want to increase the memory allocated for these components. This section describes how to increase the memory allocated for these components on systems running AIX, Linux, and Windows.

### Increasing the memory allocation for the Data server

If you run out of memory for the Data server, you might want to increase the memory allocated for the Data server. This section describes how to increase the memory allocated for the Data server on systems running AIX, Linux, and Windows.

#### Increasing the memory allocation for the Data server running on AIX:

This section describes how to increase the memory allocation for the Data server running on AIX.

To increase the memory allocated for the Data server running on AIX, complete the following steps:

1. Log in as a user with root authority.
2. To stop the Data server, run the following command:

```
stopsrc -s TSRMsr1
```

3. Run the following command:

```
odmget -qsubsysname=TSRMsr1 SRCsubsys
```

The following information is returned:

```
SRCsubsys:
 subsysname = "TSRMsr1"
 synonym = ""
 cmdargs = "-Xmx256m CLASSPATH=/usr/ibm/TPC/Data/
 cloudscape/db2j.jar:
 /usr/ibm/TPC/Data/cloudscape/db2jtools.jar LC_ALL=en_US"
 path = "/usr/ibm/TPC/Data/server/bin/aix_power/TSRMssys"
 uid = 0
 auditid = 0
 stdin = "/dev/console"
 stdout = "/dev/console"
 stderr = "/dev/console"
 action = 2
 multi = 0
```

```

contact = 3
svrkey = 0
svrmtpe = 0
priority = 20
signorm = 0
sigforce = 0
display = 1
waittime = 20
grpname = ""

```

4. Copy the entire string (excluding the quotation marks) that follows the **cmdargs** parameter.
5. Run the following command:  
`chssys -s TSRMsrv1 -a 'modified_cmdargs'`

where *modified\_cmdargs* is a modification of the string that you copied in step 4. For example, to increase the memory from 256 to 512 MB, change `-Xmx256m` to `-Xmx512m`.

6. Save the modified file.
7. Restart the Data server by running the following command:  
`startsrc -s TSRMsrv1`

#### **Increasing the memory allocation for the Data server running on UNIX or Linux:**

This section describes how to increase the memory allocation for the Data server running on Linux.

To increase the memory allocated for the Data server running on HP-UX, Linux, or Solaris, complete the following steps:

1. Log in as a user with root authority.
2. To stop the Data Server, from the command prompt, run the following command:

```
/opt/IBM/TPC/data/server/tpcdsrv1 stop
```

3. Using a text editor, open the **tpcdsrv1** file.
4. Locate the following line:

```
exec $JAVAEXE -Xrs -Xmx<XXXm> -cp $CLASSPATH com.tivoli.itsrm.server.Server &
```

where `<XXXm>` is the memory allocated for the Data server. By default, this is 256m (256 MB).

5. Increase the memory allocated for the Data server. For example, to increase the memory to 512 MB, change the line to read as follows:

```
exec $JAVAEXE -Xrs -Xmx512m -cp $CLASSPATH com.tivoli.itsrm.server.Server &
```

6. Save the modified **tpcdsrv1** file.
7. To restart the Data server, run the following command:  
`/opt/IBM/TPC/data/server/tpcdsrv1 start`

#### **Increasing memory allocation for Data server running on Windows:**

This section describes how to increase the memory allocation for the Data server running on Windows.

To increase the amount of memory available to the server, complete the following steps:

1. Select **Start** > **Run**. The Run window is displayed.

2. Type **regedit** in the Open: field and click **OK**. The Registry Editor window is displayed.
3. Expand the **HKEY\_LOCAL\_MACHINE > SOFTWARE > IBM > TSRM** folder on the left side of the Registry Editor window.
4. Right-click the **1** folder and select **New > String Value** from the pop-up menu.
5. Type **SRVJPARMS** as the name of the string.
6. Right-click the name of the string and select **Modify** from the pop-up menu.
7. Enter **-Xmx<###>m** in the Value data field, where **###** represents the number of megabytes for the server's maximum heap size. The default size is 256. Click **OK**. The largest possible value for the maximum heap size is 1536 megabytes. If the value is set to something larger than 1536, that value will be ignored and 1536 megabytes will be used as the maximum heap size.
8. Stop and restart the server to have the changes take effect. To stop the server, complete the following steps:
  - a. Go to **Start > Settings > Control Panel > Administrative Tools > Services**. The Services window is displayed.
  - b. Right-click **IBM Tivoli Storage Productivity Center - Data Server** service and click **Stop** from the pop-up menu. To restart the server, complete the following steps:
    - 1) Go to **Start > Settings > Control Panel > Administrative Tools > Services**. The Services window is displayed.
    - 2) Right-click **Tivoli Storage Productivity Center - Data Server** service and click **Start** from the pop-up menu.

### Increasing the memory allocation for the Device server

This topic describes how to increase the memory allocated for the Device server.

1. Stop the Device server.
2. Open the `server.xml` file in a text editor. This file is located at `<directory>\device\apps\was\profiles\deviceServer\config\cells\DefaultNode\nodes\DefaultNode\servers\server1`, where `<directory>` is the directory where IBM Tivoli Storage Productivity Center is installed.
3. Search for `JVMArguments`, and change the value of the `-Xmx512m` parameter.
4. Save the modified `server.xml` file.
5. Restart the Device server.

### Setting timeout values for the Device server

This topic describes how to set timeout values for the Device server.

If a probe or discovery of a storage subsystem times out before the operation completes, you receive the following error message:

```
HWN021650E Encountered timeout while connecting to CIMOM IP:port.
Check the CIMOM or increase timeout value.
```

where *IP* is the IP address, and *Port* is the port number. If you determine that the Common Information Model Object Manager (CIMOM) is not the cause of the problem, you can use the command-line interface (CLI) to increase the timeout values for the Device server.

For those storage systems that use the native interfaces (DS8000, the XIV system, SAN Volume Controller, Storwize V7000 Unified, and Storwize V7000), you see this error message:

```
HWN020103E The external process exceeded the timeout limit and was cancelled.
```

To set timeout values for the Device server, follow these steps:

1. Run the **getdscfg** command to determine the current values of the timeout properties. From the command prompt, enter the following command:

```
cli>tpctool getdscfg -user user -pwd password -url host:port
-property timeout_property timeout_value
```

where:

- *user* is an IBM Tivoli Storage Productivity Center user ID.
- *password* is the password for the Tivoli Storage Productivity Center user ID.
- *host* is the host name or IP address, and *port* is a valid port number for the HTTP service of the Device server. The default value for *port* is typically **5990**.
- *timeout\_property* is one of the following strings:
  - http.timeout
  - CIMClientWrapper.Timeout
  - Probe.Timeout.Array
  - Probe.Timeout.InBand
  - Probe.Timeout.LMM
  - Discovery.Timeout
  - CIMOMManager.TestConnectionTimeout

For the storage systems that use the native interface, the *timeout\_property* strings are:

- NAPI.Timeout.TestConnection
  - NAPI.Timeout.Probe
  - NAPI.Timeout.EventPoll
- *Timeout\_value* is the value for the timeout.

2. Run the **setdscfg** command to increase the timeout value. Run the following command:

```
cli>tpctool setdscfg -user user -pwd password -url host:port
-property timeout_property timeout_value
```

For more information about **tpctool**, see the *Tivoli Storage Productivity Center Command-Line Interface Reference*. You also can view help from the command line by issuing the command with the **-help** option.

## Changing passwords

This section provides information about changing passwords.

If you installed IBM Tivoli Storage Productivity Center using the typical installation and used the same DB2 user ID and password for the items IBM Tivoli Storage Productivity Center requires, then when you change the DB2 password, you must also change the passwords for the items that the DB2 password applies to. For example, the DB2 administrative password might also apply to the following items:

- Database administration user ID and password (for the Data Server or Device server to connect to the database)
- Database user ID and password to create the database schema
- Host authentication password (for the Storage Resource agents to communicate with the Device server)

- Storage Resource agent service login user ID and password (for Windows only, if this user ID does not exist)
- WebSphere administration user ID and password (for the Device server to communicate with embedded WebSphere if the user ID does not exist). This is the case only if you select OS authentication during the IBM Tivoli Storage Productivity Center installation. If you chose LDAP authentication, then the WebSphere administration user ID and password are set to the LDAP TPC Administrator username and password.

## Changing the DB2 user password for the Device server

This section provides information about changing the DB2 user password for the Device server.

If you change the DB2 user password, follow these steps to change it for the Device server:

1. Go to the `<TPC_install_dir>\cli` directory and run the following command:  
`tpctool encrypt <password>`

This encrypts the password.

2. Copy the encrypted password to the following file:  
`<TPC_install_dir>\device\conf\tsnmbparms.properties`

Replace this password:

```
tivoli.sanmgmt.jdbc.dbPassword=<encrypted_password>
```

3. Restart the Device server.
4. If you performed a typical installation and selected OS authentication, then this DB2 user ID is also the WebSphere administrative user ID. You must also change the WebSphere administrative user password. See “Changing the WebSphere administrative user password for the Device server.”

## Changing the WebSphere administrative user password for the Device server

This topic describes how to change the WebSphere administrative user password for the Device server.

**Note:** If you selected LDAP authentication during IBM Tivoli Storage Productivity Center installation, then the WebSphere administrative ID and password are set to the LDAP TPC Administrator username and password. Therefore, you must also change the LDAP TPC Administrator password in your LDAP-compliant directory when you change the WebSphere administrative user password for the Device server.

To change the WebSphere administrative user password for the Device server, follow these steps:

1. Stop the Device server.
2. Go to the `<TPC_install_dir>` directory. For example:  
`C:\Program Files\IBM\TPC\device\bin\W32-ix86` (Windows)  
`/<usr or opt>/IBM/TPC/device/bin/aix` (AIX)  
`/<usr or opt>/IBM/TPC/device/bin/linux` (Linux)
3. Run the **ChangeWASAdminPass** script:  
`ChangeWASAdminPass.<bat or sh> <user_ID> <password> <install_dir>`

Where `<user_ID>` is the WebSphere user ID, `<password>` is the password, and `<install_dir>` is the directory where the Device server is installed.

4. Restart the Device server.

## Changing the DB2 password for the Data Server

This section provides information about changing the DB2 password for the Data Server.

To change the DB2 password, complete the following steps:

1. From a command prompt, change to the directory where the file `repository.config` resides in the IBM Tivoli Storage Productivity Center installation path. For example:  

```
c:\Program Files\IBM\TPC\data\config\ (Windows)
/<usr or opt>/IBM/TPC/data/config/ (AIX or Linux)
```
2. Run the following command for the Data server. This assumes that the Tivoli Storage Productivity Center directory is in `C:\Program Files\IBM\TPC` for Windows.  

```
"C:\Program Files\IBM\TPC\data\config>"program files"\
ibm\tpc\jre\bin\java -classpath "\program files\ibm\tpc\data\server\lib\
TSRMSrv.zip" com.tivoli.itsrm.repository.Transform
-p <new_password> repository.config
```
3. Stop and then restart the Data Server.
4. Start the Tivoli Storage Productivity Center GUI.

## Changing the host authentication password for the Device server

This topic describes how to change the host authentication password for the Device server. If you change this password for the Device server, you must also change the host authentication password for the Fabric agent.

To change the host authentication password for the Device server, complete the following steps:

1. Open the IBM Tivoli Storage Productivity Center GUI.
2. Expand **Administrative Services > Services > Device server > Device server**.
3. Right-click **Device server** and select **Change password** from the pop-up menu.
4. Enter the new password and click **OK**.
5. Make sure you change the host authentication password for the Fabric agent. See "Changing the host authentication password for the Fabric agent."

## Changing the host authentication password for the Fabric agent

If you change the host authentication password for the Device server, you must also change the password for the Fabric agent. This topic describes how to change the host authentication password for the Fabric agent.

To change the host authentication password for the Fabric agent, complete the following steps:

1. For Windows, go to the following directory and run these commands:  

```
cd TPC_installation_directory\ca\subagents\TPC\Fabric\bin\w32-ix86
setenv
srmcp -u user_ID -p password ConfigService
setAuthenticationPW new_host_password
```
2. For UNIX or Linux, go to the following directory and run these commands:  

```
cd TPC_installation_directory/ca/subagents/TPC/Fabric/bin/
aix_or_linux_or_solaris2
. ./setenv.sh (note there is a space between the two dots)
./srmcp.sh -u user_ID -p password ConfigService
setAuthenticationPw new_host_password
```

## Using the password tool

Use this procedure to change the password for the DB2 user ID, DB2 services, and Tivoli Storage Productivity Center for Replication.

The tool makes modifications to the DB2, Tivoli Storage Productivity Center for Replication, and WebSphere applications so that they can authenticate properly to the operating system whenever you change the administrator password.

Follow these steps to change the DB2 user ID, DB2 services, and Tivoli Storage Productivity Center for Replication password.

**Important:** Use this procedure only if the administrator ID matches the DB2 user ID, DB2 services, and Tivoli Storage Productivity Center for Replication password.

1. Log on to Tivoli Storage Productivity Center for using the administrator user ID.
2. Change the password on the system.

### For Windows

Follow these steps:

- a. From the desktop, click **Start > Control Panel**.
- b. On the Control Panel window, double-click **User Accounts**.
- c. Click **Change your password**.
- d. In the Set Password window, enter the new password twice. Click **Change password** to change the password. The User Accounts window is displayed again. Close the window.

### For AIX or Linux

Run the **passwd** command logged in with the DB2 user ID. To change other user credentials, run the **passwd username** command.

3. Open a command prompt window and change the directory to the following default directory:

```
C:\Program Files\IBM\TPC\data\server\tools (for Windows)
/opt/IBM/TPC/data/server/tools (for Linux or UNIX)
```

Run the following program to start the password tool.

```
changepasswords.bat (for Windows)
changepasswords.sh (for Linux or AIX)
```

4. The Choose the password to change window is displayed. Select the **Change DB2 user password for Device and Data Server** option and click **OK**.
5. The Change DB2 user password for Device and Data Server window is displayed.
6. Perform the following steps:
  - a. Type the password for the DB2 user ID in both the **Enter password** and **Confirm password** fields.
  - b. The **Restart the Data and Device Server** option is selected. Leave it as is and click **OK**.
  - c. When you see the following message, click **Yes**.  
Are you sure you want to change the password?
  - d. The Changing DB2 password for user Administrator window is displayed and shows the progress messages while the password is changed. The process takes approximately 5 minutes.

**Note:** The user "Administrator window" can be "db2admin window" or "db2inst1 window."

- e. When the message Complete is displayed, click **Back to Main**.
- f. The Choose the password to change window is displayed again.

**Note:** For UNIX, the password tool will not check if the WebSphere administrative user ID is the same as the DB2 user ID. The tool will not change the password. In a typical installation, the WebSphere administrative user ID and the DB2 user ID are the same and the WebSphere administrative password must be changed before running the password tool. If this step is not done, the password tool will display an error message when restarting the Device server. If this action occurs, update the WebSphere administrative password and start the Device server again.

7. For Windows only: Select **Change password for the DB2 services** and click **OK**.
8. For Windows only: The Changing DB2 Services password window is displayed. Follow these steps:
  - a. Type the user name of the DB2 Services in the **DB2 Services Username** field.
  - b. Type the password that is the same as the password for the user of the DB2 Service in both the **Enter password** and **Confirm password** fields. Click **OK**.
  - c. When you see the following message, click **Yes**:  
Are you sure you want to change the password?
  - d. The Changing DB2 Services password window is displayed and shows the progress messages while the password is changed.
  - e. When the message Complete is displayed, click **Back to Main**.
9. The Choose the password to change window is displayed again. Select **Change TPC-R password** to change the password for Tivoli Storage Productivity Center for Replication and click **OK**.
10. The Change TPC-R password window is displayed. Follow these steps:
  - a. Type the password that is the same as the DB2 user ID password in both the **Enter password** and **Confirm password** fields. Click **OK**.
  - b. When you see the following message, click **Yes**:  
Are you sure you want to change the password?
  - c. The Changing TPC-R password window is displayed and shows the progress messages while the password is changed.
  - d. When you see the message Complete, click **Back to Main**.
  - e. The Choose the password to change window is displayed again.

**Note:** To verify the password changes has completed successfully, check the log located in *TPC\_install\_directory\Data\log\PWTool.log*.

11. Click **Exit program** to exit the password tool.
12. Press any key to continue, and close the command prompt window. This completes the password change.

**Note:** On Windows, if the WebSphere administrative user ID for the Device server is the same as the DB2 user ID, the password tool also changes the WebSphere administrative password.



## Checking for a fully qualified host name

IBM Tivoli Storage Productivity Center requires fully qualified host names. Some machines might be configured to return a short host name, such as **ddunham** instead of a fully qualified host name, such as **ddunham.myorg.mycompany.com**. This topic provides information on how to check for a fully qualified host name.

### Checking for a fully qualified host name for Windows systems

This topic provides information on how to verify a fully qualified host name for Windows systems.

To verify a fully qualified host name, complete the following steps:

1. On the desktop, right-click **My Computer**.
2. Click **Properties**.
3. The System Properties panel is displayed.
4. On the Computer Name tab, click **Change**.
5. The Computer Name Changes panel is displayed. Enter your fully qualified host name in the Computer name field. Click **More**.
6. The DNS Suffix and NetBIOS Computer Name panel is displayed. Verify that the **Primary DNS suffix** field displays a domain name. Click **OK**.
7. Click **OK** on the Computer Name Changes panel.
8. Click **Apply** and close the System Properties panel.

### Checking for a fully qualified host name for AIX systems

This topic provides information on how to verify a fully qualified host name for AIX.

The default domain name search order is as follows:

1. Domain Name System (DNS) server
2. Network Information Service (NIS)
3. Local `/etc/hosts` file.

If the `/etc/resolv.conf` file does not exist, the `/etc/hosts` file is used. If only the `/etc/hosts` file is used, the fully qualified computer name must be the first one that is listed after the IP address.

Verify that the `/etc/resolv.conf` file exists and contains the appropriate information, such as:

```
domain mydivision.mycompany.com
nameserver 123.123.123.123
```

If NIS is installed, the `/etc/irs.conf` file overrides the system default. It contains the following information:

```
hosts = bind,local
```

The `/etc/netsvc.conf` file, if it exists, overrides the `/etc/irs.conf` file and the system default. It contains the following information:

```
hosts = bind,local
```

If the `NSORDER` environment variable is set, it overrides all of the preceding files. It contains the following information:

```
export NSORDER=bind,local
```

## Checking for a fully qualified host name for Linux systems

This topic provides information on how to verify a fully qualified host name for Linux.

Linux uses a resolver library to obtain the IP address corresponding to a host name. The `/etc/host.conf` file specifies how names are resolved. The entries in the `/etc/host.conf` file tell the resolver library what services to use, and in what order, to resolve names. Edit the `host.conf` file using the `vi` editor to add the following lines:

```
Lookup names through DNS first then fall back to /etc/hosts.
order bind,hosts
Machines with multiple IP addresses.
multi on
Check for IP address spoofing.
nospoof on
```

The **order** option indicates the order of services. The sample entry specifies that the resolver library should first consult the name server to resolve a name and then check the `/etc/hosts` file. It is recommended to set the resolver library to first check the name server, `bind` file, and then the `hosts` file (`hosts`) for better performance and security on all your servers. You must have the DNS and BIND software installed for this configuration to work.

The **multi** option determines whether a host in the `/etc/hosts` file can have multiple IP addresses. Hosts that have more than one IP address are said to be multihomed, because the presence of multiple IP addresses implies that the host has several network interfaces.

The **nospoof** option takes care of not permitting spoofing on this machine. IP-Spoofing is a security exploit that works by tricking computers into a trust relationship that you are someone that you really are not. In this type of attack, a machine is set up to look like a legitimate server and then issue connections and other types of network activities to legitimize end systems, other servers, or large data repository systems. This option must be set ON for all types of servers.

## Checking for a fully qualified host name for Oracle Solaris

This topic provides information about how to verify a fully qualified host name for Oracle Solaris systems.

Verify that the `/etc/resolv.conf` file exists and contains the appropriate information, such as:

```
domain mydivision.mycompany.com
nameserver 123.123.123.123
```

A short name is used if the `/etc/nsswitch.conf` file contains a line that begins as follows and if the `/etc/hosts` file contains the short name for the computer:

```
hosts: files
```

To correct this problem, follow these steps:

1. Change the line in the `/etc/nsswitch.conf` file to the following:  

```
hosts: dns nis files
```
2. Enter the following command to stop the **inet** service:  

```
/etc/init.d/inetsvc stop
```
3. Enter the following command to restart the **inet** service:  

```
/etc/init.d/inetsvc start
```

## Granting local administrative privileges to a domain account

For Windows users, the user account which the Storage Resource agent runs under requires local administrative rights. Because these rights are not necessarily guaranteed for domain users in a Windows domain environment, this topic provides information about how to grant local administrative rights for domain users. Using this procedure, you do not have to manually process each machine in the domain.

**Note:** These steps are for a Windows system that is a member of a Windows domain and not for the Windows Domain Primary Domain Controller.

To use Group Policy to grant local administrative privileges to a domain account, complete the following steps:

1. On the domain controller, go to **Administrative Tools > Active Directory Users and Computers** (you must be running with Domain Administrator privileges).
2. Right-click the Organizational Unit (OU) upon which you want to apply the Group Policy. Click **Properties**.
3. The Group Policy Properties panel is displayed. Select the Group Policy tab and click **New** to create a Group Policy.
4. Designate a name for the new Group Policy. Select the new Group Policy and click **Edit**.
5. The Group Policy Object Editor panel is displayed. Go to **New Group Policy Object *your\_policy* > Computer Configuration > Windows Settings > Security Settings > Restricted Groups**. Right-click **Restricted Groups**. Click **Add Group**.
6. For example, name the new group "Administrators." Under "Properties", add the user "Administrator", and the domain accounts or groups upon which you want the Group Policy in effect for. For example, you can add "TPC\tapeadmin", "TPC\tapegroup", and "TPC\TestGroup". Click **OK**.
7. Add these user rights to the domain account:
  - Act as part of the operating system
  - Log on as a service

In the Group Policy Object Editor, go to **New Group Policy Object *your\_policy* > Computer Configuration > Windows Settings > Security Settings > Local Policies > User Rights Assignments**. In the content pane, select "Log on as a service" and double-click. Add the domain user for whom you are granting the user right for and click **OK**. Repeat this step for "Act as part of the operating system."

8. The group policy is now enforced for the Organizational Unit to include the domain accounts and groups specified under the local Administrators group on each computer in the Organizational Unit. In addition, the domain user has been granted the necessary rights. To verify the user rights, log in to a domain computer and open the Computer Management console. Select **Groups**, double-click the Administrators group, and verify the membership of the domain users.

## Importing SSL certificates for VMware

If your VMware uses SSL certificates, you must import the SSL certificates for VMware Virtual Infrastructure Data Source.

If your VMware uses SSL certificates, you must follow these steps before you can get data from the VMware Virtual Infrastructure. Complete the following steps:

1. Open a browser. Point the browser to the following:  
`https://ESX_server_IP_address/host`
2. You will see a list of Configuration files. Right-click `ssl_cert`. Save the file as `ESX_machine_name.crt`.

**Note:** Depending on your browser, you would save the file using **Save Target as** (for Internet Explorer) or **Save Link as** (for Firefox).

Use this file to import the certificate.

3. The VirtualCenter Server uses a self-generated SSL certificate located in the following directory:

```
C:\Documents and Settings\All Users\Application Data\
VMware\VMware VirtualCenter\SSL\rui.crt
```

Install these certificates into a certificate store on the client machine. Copy the certificate file or files from the servers to a directory on the client machine.

Enter the **keytool** command:

```
TPC_installation_directory/jre/bin/keytool -import
-file certificate-filename
-alias server-name
-keystore vmware.jks
```

For example, run this command:

```
TPC_installation_directory/jre/bin/keytool -import
-file vc-server1.crt -alias vc-server1
-keystore vmware.jks
```

Enter a password for the keystore. Enter **yes** to import the certificate.

The truststore will be located in the Device server configuration directory:

```
TPC_installation_directory/device/conf
```

If the truststore was generated in another directory, it must be copied to `TPC_installation_directory/device/conf`. The truststore is called **vmware.jks**.

The truststore will automatically be defined at service startup time as this property in the Device server JVM:

```
javax.net.ssl.trustStore System
```

4. Add a VMware data source using the URL, user name, and password for the service console.
5. Run a discovery job for the VMware data source.
6. Run a probe job for the VMware data source.
7. You will now be able to display VMware Virtual Infrastructure in reports and the topology viewer.

**Tip:** In a Linux operating system environment, if you install Java when installing the operating system, the operating system version of **keytool** is installed. The operating system level of the **keytool** might not be the correct version for Tivoli Storage Productivity Center to use. To create the **vmware.jks** file correctly, use the **keytool** located in `/opt/IBM/TPC/jre/bin`.

## SQL access to Tivoli Storage Productivity Center views

IBM Tivoli Storage Productivity Center stores the information collected by data collection jobs in its DB2 database repository. This information is organized into a set of accessible views. You can use Structured Query Language (SQL) to retrieve the information from these views to create custom reports or for use within other applications.

**Prerequisite:** See the *Planning* section in the Tivoli Storage Productivity Center Information Center for information about the requirements when retrieving data from views in the database repository.

A view is a logical table that consists of data that is generated by a query. A view is based on an underlying set of base tables in the database repository, and the data in a view is determined by a SELECT statement that is run on those base tables. The data collected Tivoli Storage Productivity Center is organized into the following sets of views that you can access in the database repository:

### **Storage entity views**

These views include detailed information about the properties of storage entities. For example, the name, disk space, and available disk space for a storage subsystem. Information about the following storage entities are available in views in the database repository:

- Storage entity views: Hosts, volume groups, disks, virtual disks, external storage volumes, disk paths, virtual disk paths, controllers, logical volumes, file systems, clusters, cluster resource groups, hypervisors, virtual machines, virtual machine disks, hypervisor volumes, host virtual files, hypervisor partitions, MAC addresses, files, file types, directories, exports, OS users, OS user groups, IP addresses, Windows domains, and Storage Resource agents.
- Host relationship views:
  - Logical volume to host disk
  - Logical volume to logical volume (Parent and child logical volume relationships)
  - Computer to storage entity (disk or logical volume)
  - Filesystem to logical volume(for example, the relationship between a logical volume and a host disk)
- Database views: Database instances, database containers, database archive logs, database control files, database datafiles, database devices, database log files, database log groups, database segments, database tables, database tablespaces
- Storage subsystem views: Storage subsystems, storage virtualizers, storage pools, extents, volumes, disk drives, disk groups, element managers, nodes, backend disks, storage virtualizer I/O groups, volume paths, array sites, ranks
- Storage subsystem relationship views:
  - Disk drive to storage extent
  - Storage volume to storage extent
  - Volume to disk drive
  - Storage volume to backend disk
- NAS Filers
- Tape library views: Tape libraries, tape cartridges, tape drives, tape frames, tape I/O ports, tape media changers, tape media locations
- Fabric views: Fabrics, nodes, peripheral entities, port to port connections, switches, switch blades, zones, zone members, zone sets, zone aliases, FC ports, host bus adapters
- Fabric relationship views:
  - Blade to FC port
  - Node to HBA, Node to port

- Peripheral entity to node
- Switch to port
- Zone to alias, Zone to zone member, Zone to node, Zone to member
- Zone set to zone
- Zone alias to node
- Alias to member
- Fabric to switch, Fabric to zoneset, Fabric to zone, Fabric to zone alias
- Performance management views (for example, metrics and threshold violations)

### **Entities defined by Tivoli Storage Productivity Center**

These views include information about entities that are defined by Tivoli Storage Productivity Center. These entities include Data agents, Fabric agents, the alert log, the Tivoli Storage Productivity Center servers, scan profiles, external tools, monitoring groups, reporting groups, and storage resource groups.

### **Data sources**

These views include information about the data sources that are associated with Tivoli Storage Productivity Center. These data sources include Data agents, Storage Resource agents, Fabric agents, Tivoli Storage Productivity Center servers, CIMOM agents, Out of Band agents, and VMWare Virtual Infrastructure.

### **Aggregated views**

These views provide summary information for the database history, data in a database instance, and the Data agent file system.

### **Reporting views**

These views combine several different entities into one view, such as aggregate database and file system summary views, file size, time, and age distribution views, enterprise-wide statistics views, directory size views, and historical views of statistics for database summaries, ping jobs, user space usage, and files in user groups.

### **Rollup views**

These views include rollup report information from the master and subordinate Tivoli Storage Productivity Center servers, Data agents, Fabric agents, Storage Resource agents, host cluster data, computer group, host, database computer groups, fabric SAN assets, switch assets, storage subsystem group, storage subsystems, and Data Manager for Databases.

See the Tivoli Storage Productivity Center's support Web site for detailed descriptions of the views provided with Tivoli Storage Productivity Center: [http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli\\_Storage\\_Productivity\\_Center\\_Standard\\_Edition](http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center_Standard_Edition).

This Web page contains the following additional documentation:

- Information about the TPCREPORT schema, including:
  - A list of the views that exist in the schema and descriptions of the columns with the views.
  - A list of the views that contain performance data and descriptions of the performance metrics within those views.
  - Diagrams that show the relationships between views.

- A diagram that traces the relationship between a computer and a storage subsystem disk and shows the views that contain information about the storage entities within that relationship.

### Accessing views in the database repository

Use SQL to access the views in Tivoli Storage Productivity Center database repository. You can do this using an SQL command line or an application capable of using SQL to retrieve information from views. This topic provides an example of how to create and view a report based on those views.

To retrieve information from the database repository you must have the proper authority to access the views. You must also collect information about your storage assets to populate those views with information. See the *Planning* section in the Tivoli Storage Productivity Center Information Center for more information.

To perform the steps described in this topic, you must install the correct version BIRT. You can download BIRT v2.2.1 from the following Web page: <http://www-01.ibm.com/software/brandcatalog/portal/opal/details?catalog.label=1TW10OT02>. You must use this version of BIRT when working with Tivoli Storage Productivity Center and Tivoli Common Reporting.

You can use a number of SQL and reporting tools for accessing and viewing the content of IBM Tivoli Storage Productivity Center's exposed views. For example, you can use the DB2 Command Line Processor to run SQL commands against the views.

This topic describes how to use BIRT to create a report package based on the views and how to import that package into Tivoli Common Reporting for viewing. Tivoli Common Reporting is installed with Tivoli Storage Productivity Center.

See the following web sites for more information about BIRT and Tivoli Common Reporting:

- BIRT overview information: <http://www.eclipse.org/birt/phoenix/intro/>
- Tivoli Common Reporting overview information: <http://www.ibm.com/developerworks/spaces/tcr>

Complete these steps to use BIRT and Tivoli Common Reporting to work with Tivoli Storage Productivity Center's views:

1. Install BIRT on a computer that has access to Tivoli Storage Productivity Center's database repository. The database repository is located on the computer where the Tivoli Storage Productivity Center server is installed.
2. Start the BIRT Designer tool.
3. Set the connection information for the Tivoli Storage Productivity Center database repository in Data Source section of the user interface.
4. Use the **Data Sets** section of the user interface to create a report based on SQL queries to the exposed views that are in the TPCREPORT schema. See <http://www.eclipse.org/birt/phoenix/intro/> for information on how to use BIRT to create a report. When writing queries against the views make sure to qualify the view names with the TPCREPORT schema (for example, TPCREPORT.STORAGESUBSYSTEM).

**Note:** BIRT limits the amount of views that are displayed to 100 by default. To change this, access the **Preferences** window for the Data Set Editor and enter a value greater than 275 in the **Maximum number of tables in each schema to display** field.

5. Define the layout of the report. Use BIRT to customize the appearance of a report, including customizing column names and creating charts, graphics, and aggregated totals for the data in the report. Use the preview function to view a snapshot of the report that you defined.
6. Create a report package (archive file) for the reports to import into Tivoli Common Reporting.
7. Import the report package into Tivoli Common Reporting.
  - a. Start Tivoli Integrated Portal. For information about starting Tivoli Integrated Portal, go to the Tivoli Storage Productivity Center information center. Search for **Starting the Tivoli Storage Productivity Center GUI from Tivoli Integrated Portal**.
  - b. Expand **Reporting** in the navigation tree and click **Common Reporting**.
  - c. Right-click **Report Sets** and select **Import Report Package**.
  - d. Select the report package you want to import and click **Import**. Click **Overwrite** in the **Advanced Options** section if the report package was previously imported. The name of the report package appears in the **Reports** section of the panel.
  - e. Right-click the report package and select **Views As > output\_type**, where *output\_type* represents the format of the generated report.
  - f. View the report.

## Repository Copy tool

You can use the Repository Copy tool to export all the tables in the IBM Tivoli Storage Productivity Center repository (Data Manager data) for purposes of debugging problems.

You can send the exported data to the IBM Support Center to help debug problems.

### Note:

- You must stop the Device server and Data server before running **repocopy**.
- You cannot import **repocopy** data into the IBM Tivoli Storage Productivity Center database if the migration tool has not been run against the database. When you upgrade the Device server, you must run the migration tool before importing **repocopy** data.

## Exporting repository data

Use the repocopy tool to export data from an existing repository into a comma delimited file.

To export repository data, follow these steps:

1. Go to the following default directory:
  - If on Windows: c:\Program Files\IBM\TPC\data\server\tools
  - If on Linux or UNIX: /opt/IBM/TPC/data/server/tools
2. Run the **repocopy.bat** or **repocopy.sh** command. A window is displayed prompting you for what you would like to do.
3. Select **Export data from repository tables**, and then click **Next**. The Options for Import/Export window is displayed.
4. In the Options for Import/Export window, enter information in the following fields:



**Directory for Export**

Enter the directory where the comma delimited file is saved.

**Delimiter**

Enter a delimiter for the delimited file format (comma is the default).

**Quote** Enter the symbol that contains string data (double quotation marks is the default).

IBM Tivoli Storage Productivity Center exports the data into the comma delimited file you specify, and place it in a file named *tablename.txt*. Click **Next**.

5. The Choices for DB2 Export window are displayed. You have the following choices:

- Export Using DB2 native format
- Export Using text files

Make a choice and click **Next**.

6. The Table Export Choices windows is displayed. You have the following choices:

- Export base tables (always)
- Export Performance Manager tables
- Export history tables used for IBM Tivoli Storage Productivity Center for Data history reports

Make a choice and click **Next**.

7. The Connection Properties window is displayed.

In this step, Tivoli Storage Productivity Center reads the `server.config` file and determines your current repository and the information for accessing the database.

The information detected in the `server.config` file is displayed in the Connection Properties window within the following fields: **Database Types**, **User name**, **Password**, **Driver Class**, **Driver URL**, **Database**, **DB Creator**, **Classpath**.

**Note:** If you want to export data from a different database from the one listed in the `server.config` file, you can select the database from the **Database Types** list box. Manually enter the database information. Click **Finish**.

8. Tivoli Storage Productivity Center then connects to the database and displays the database and connection information. Click **Run** to begin the export process.
9. A window is displayed containing the export progress log. As you progress through the export process, messages are written to this progress log. You can keep track of the steps through the progress log.

**Note:** When you run the **repocopy** tool on a server that has a remote database, the **repocopy** tool displays a message. When the **repocopy** tool is used with a remote database, the DB2 shared library is unavailable for loading `libTSRMinsudb.so`. You can ignore this message. Click **OK** and continue.

**Importing repository data**

Use the Repository Copy tool to import data from a comma delimited file. You must have previously exported data from an existing repository.

To import data into repository tables, follow these steps:

1. Stop the Device server and Data server.
2. Go to the following directory:
  - If on Windows: c:\Program Files\IBM\TPC\data\server\tools
  - If on Linux or UNIX: /opt/IBM/TPC/data/server/tools
3. Run the **repocopy.bat** or **repocopy.sh** command. A window is displayed prompting you for what you would like to do
4. Select **Import data into repository tables**, and then click **Next**. The Options for Import/Export window is displayed.
5. In the Options for Import/Export window, enter information in the following fields:

**Directory for Import**

Enter the directory where the comma delimited files are stored.

**Delimiter**

Enter a delimiter used for the delimited file format (comma is the default). You must use the same delimiter that was used in the exported file.

**Quote** Enter the symbol used to contain string data (double quotation marks is the default). You must use the same double quotation marks that were used in the exported file.

**Delete before inserting**

Check this option if you want to delete any existing data in the repository tables before importing new data.

IBM Tivoli Storage Productivity Center imports the data from the comma delimited file that you specify. Click **Next**.

6. The Connection Properties window is displayed.
7. Enter the database and the access information of the database to which you want to import data. The Connection Properties window contains the following fields: **Database Types**, **User name**, **Password**, **Driver Class**, **Driver URL**, **Database**, **DB Creator**, **Classpath**. Click **Finish**.
8. Tivoli Storage Productivity Center then connects to the database and displays the database and connection information. Click **Run** to begin the import process.
9. A window is displayed containing the import progress log. As you progress through the import, messages are written to this progress log. You can use the progress log to track the steps as they happen.

**Note:** When you run the **repocopy** tool on a server that has a remote database, the **repocopy** tool displays a message. When the **repocopy** tool is used with a remote database, the DB2 shared library is unavailable for loading libTSRMinsudb.so. You can ignore this message. Click **OK** and continue.

## Service tool: collecting information

The Service tool collects information from all installed IBM Tivoli Storage Productivity Center components. The tool detects system configuration, collects the applicable information, and creates a compressed file that can be sent to the IBM support center.

The Service tool collects the following information:

- Host name
- IP address and configuration information

- Operating system and version
- Java™ home, version, and class path
- Java Virtual Machine (JVM) implementation name and version
- Protocol statistics and current IBM Tivoli Storage Productivity Center / Internet Protocol network connections, including listening ports
- Diagnostic information regarding the system and its services
- Listing of all library files, for example, server and library, agent and library, and GUI library

When the Service tool is run on the system where the Data server or the Device server is installed, it also collects the following information:

- For the Data server, information about all of the graphical user interfaces (remote and local) that are associated with this Data server
- All applied e-fixes
- Installation logs
- Everything in the log directory, including subdirectories
- Everything in the conf directory
- Directory listing of the lib directory
- Everything in the log and conf directories underneath the Websphere directory
- The version of DB2 (information from the db2level command)
- Information from the **ipconfig /all** command (Windows operating systems)
- Information from the **ipconfig -a** command (Linux and UNIX operating systems)
- Information from the **netstat -an** command

When the Service tool is run on the system where the database repository is installed, it also collects the following information:

- The DB2 support information (information from the db2support command)

When the Service tool is run on the system where Tivoli Integrated Portal and the TCR component is installed, it also collects the following information:

- Everything in the log and conf directories
- Logs for the deployment engine
- Listing of COI components installed
- InstallAnywhere logs for Tivoli Integrated Portal and TCR
- TCR COI related logs
- Tivoli Integrated Portal Websphere configuration files

When the Service tool runs on the Storage Resource Agent computer, it collects the following information:

- All applied e-fixes
- Everything in the log directory (including subdirectories).
- Everything in the conf directory
- Directory listing of the lib directory
- Information from the **ipconfig /all** command (for Windows operating systems).
- Information from the **ifconfig -a** command (for UNIX and Linux operating systems).
- Information from the **netstat -an** command

By default, the service data is collected in the following directory:

**On Windows operating systems:**

*TPC\_installation\_directory\service\data*

**On UNIX and Linux operating systems:**

*TPC\_installation\_directory/service/data*

You can change the default directory.

**Tip:** You can run the Service tool on IBM Tivoli Storage Productivity Center regardless of whether you configured it on a single server or on a split server. The Service tool automatically recognizes the installed components and collects service data about them. To learn how to run the Service tool for servers, see the following topic: "Creating a compressed file about servers for the IBM Support Center."

### **Creating a compressed file about servers for the IBM Support Center**

Run the Service tool for the Data server and the Device server to create a compressed file that can be sent to the IBM Support Center.

You must have administrator authority on Windows operating systems or root authority on UNIX and Linux operating systems.

To run the tool for the servers, follow these steps:

1. Log on to the system.
2. If you used the default directory, go to the following directory:
  - On Windows operating systems: C:\Program Files\IBM\TPC\service\.
  - On Linux or UNIX operating systems: */usr\_or\_opt/IBM/TPC/service/*.
3. Run the following program:
  - On Windows operating systems: *service.bat*.
  - On Linux or UNIX operating systems: *service.sh*.

A compressed file, *TPCServiceFiles.zip*, is created in the following directory:

**On Windows operating systems**

*TPC\_installation\_directory\service\data*

**On Linux or UNIX operating systems**

*TPC\_installation\_directory/service/data*

**Related tasks:**

"Creating a compressed file about Storage Resource Agents for the IBM Support Center"

Run the Service tool on Storage Resource Agents that were deployed by using the graphical user interface to create a compressed file that can be sent to the IBM Support Center.

### **Creating a compressed file about Storage Resource Agents for the IBM Support Center**

Run the Service tool on Storage Resource Agents that were deployed by using the graphical user interface to create a compressed file that can be sent to the IBM Support Center.

You must have administrator authority on Windows operating systems or root authority on UNIX and Linux operating systems.

To run the Service tool on the Storage Resource agents, do the following steps:

1. Log on to the system.
2. In the navigation tree, expand **Administrative Services > Data Sources > Data/Storage Resource Agents**.
3. In the content pane, select the agent. From the **Select Action** menu, click **Collect Service Data**.

The following compressed files are created:

**On Windows operating systems:**

C:\Program Files\IBM\TPC\data\log\SRATraces\*SRA\_computer\_name*\TPCServiceInfo.zip

**On UNIX and Linux operating systems:**

/opt/IBM/TPC/data/log/SRATraces/*SRA\_computer\_name*/TPCServiceInfo.zip

where *SRA\_computer\_name* represents the name of the computer on which the Storage Resource agent is located.

If the compressed file cannot be created, a message indicates that the job was unsuccessful.

**Related tasks:**

“Creating a compressed file about servers for the IBM Support Center” on page 222

Run the Service tool for the Data server and the Device server to create a compressed file that can be sent to the IBM Support Center.

## Customizing the service tool from the command line

Using various parameters, you can use the service tool to collect data about specific IBM Tivoli Storage Productivity Center components. You can also use the command line parameters to specify a location to place the data that is collected, specify whether the data should be compressed, or both.

### Help and output command-line parameters

To customize the data that is collected by the service tool, use the following command-line parameters:

**-output** *directory\_path*

Places the files that contain the data that was collected in a directory that you specify. If you specify a directory that does not exist on your system, that directory is created. If you do not use the **-output** *directory\_path* parameter, the files are placed in the default directory:

**On Windows operating systems:**

*TPC\_installation\_directory*\service\data

**On Linux or UNIX operating systems:**

*TPC\_installation\_directory*/service/data

**Restriction:** If you specify a directory, the directory path cannot contain spaces.

**-nozip**

When you use this parameter, separate directories are created for each Tivoli Storage Productivity Center component for which data was collected. You can then create compressed files for the files. In this way, you can control the size

and content of each compressed file. To specify a directory other than the default directory, use the `-output directory_path` parameter.

**Tip:** You can specify more than one parameter, for example, `C:\Program Files\IBM\TPC\service\service.bat -device -nozip`.

## Collecting data for specific Tivoli Storage Productivity Center components

You can use the service tool to collect data about specific Tivoli Storage Productivity Center components.

The following list shows the parameters that are used to collect data for each component:

**-all**

Collects data about all Tivoli Storage Productivity Center components.

**-data**

Collects data about the data server component files.

**-device**

Collects data about the device server component files.

**-repl**

Collects data about the replication server component files.

**-sra**

Collects data about the Storage Resource Agent component files.

**-cli**

Collects data about command-line interface files.

**-gui**

Collects data about IBM Tivoli Storage Productivity Center graphical user interface files.

**-tip**

Collects data about Tivoli Integrated Portal interface files.

**-filefilter *name***

Applies a filter when you search for files. The following examples show ways in which you can filter data:

- To find all of the files and folders that contain the string *ver* in their names, issue the following command: `>./service.sh -filefilter ver`
- To find all of the files and folders within the data or graphical user interface components that contain the string *ver* in their names, issue the following command: `>./service.sh -data -gui -filefilter ver`

When you use the **-filefilter** parameter, the resulting compressed file contains archive directories for every component that you specified. The content of these archive directories is filtered and contains only files and folders in which the string *ver* is used, as shown in the previous example.

**Restriction:** The **-filefilter** parameter does not apply a filter when you collect data about Storage Resource agents, which are specified by using the **-sra** parameter.

**Tip:** You can specify more than one parameter, for example, `C:\Program Files\IBM\TPC\service\service.bat -cli -gui`.

**Attention:** An existing compressed file is overwritten when another file of the same name is created. For example, if you run `C:\Program Files\IBM\TPC\service\service.bat -sra -gui`, a file named `TPCServiceFiles.zip` is created. If you rerun the tool with the same component options, a new file named `TPCServiceFiles.zip` is created. This new file overwrites the previously created file unless you specify the **-nozip** parameter, or use the **-output** parameter to specify a different path.

If you specify the **-nozip** parameter, a subdirectory is created for the components that you specified in the parameter. If you did not specify a specific component, data is collected for all installed components, and the data is placed in files in the following directory:

**Windows operating systems:**

*TPC\_installation\_directory\service\data*

**UNIX or Linux operating systems:**

*TPC\_installation\_directory/service/data*

Files for particular components are then placed in a subdirectory that corresponds to that component. Common files, such as `license.txt`, are placed in the following directory:

**Windows operating systems:**

*TPC\_installation\_directory\service\data*

**UNIX or Linux operating systems:**

*TPC\_installation\_directory/service/data*

---

## Administering Tivoli Integrated Portal

This topic provides information on how to perform administrative tasks in IBM Tivoli Integrated Portal that are reflected in IBM Tivoli Storage Productivity Center.

### Changing the user authentication method

The IBM Tivoli Storage Productivity Center installation program enables you to select a user authentication method that is used by Tivoli Storage Productivity Center, IBM Tivoli Storage Productivity Center for Replication, and IBM Tivoli Integrated Portal. You can choose to authenticate users against the users defined for the local operating system or to authenticate users against the users defined in a Lightweight Directory Access Protocol (LDAP) or Microsoft Active Directory repository. You can change the user authentication method after installation using Tivoli Integrated Portal.

You can select the following user authentication methods in Tivoli Integrated Portal:

- **Federated repositories.** This method authenticates the Tivoli Storage Productivity Center user against a Lightweight Directory Access Protocol (LDAP) or Microsoft Active Directory, directory service.
- **Local operating system.** This authentication method authenticates the Tivoli Storage Productivity Center user against the users defined for the local operating system. If you use operating system authentication, the use of the single sign-on feature is limited. OS authentication does not support single sign-on for element managers, even when the element manager is installed on the same computer as Tivoli Storage Productivity Center.

Whenever you change the user authentication method (from local operating system to LDAP or from LDAP to local operating system), you must first back up the Device server and Tivoli Integrated Portal WebSphere configuration files. The WebSphere configuration backup files provide a restoration point in the event of a failure when performing these procedures. You must also back up the WebSphere configuration files after you have changed the user authentication method to LDAP and want to later change the LDAP user authentication.

When you install Tivoli Integrated Portal, the OS user name and password are used to create a file user registry with the same user name and password. When you change from OS to LDAP authentication, the file user registry is kept in the Tivoli Integrated Portal registry. This step is done to always allow for a valid user ID to log in and configure Tivoli Integrated Portal, even if the LDAP server is down.

The following topics describe how to select and configure each user authentication method.

### **Changing the authentication method from local operating system to LDAP federated repositories**

You can configure IBM Tivoli Storage Productivity Center to communicate with an external Lightweight Directory Access Protocol (LDAP) repository such as IBM Tivoli Directory Server or Microsoft Active Directory. To change the authentication method from local operating system to LDAP, you must use IBM Tivoli Integrated Portal.

To change the user authentication method from local operating system to LDAP, complete the following steps. Contact your LDAP server administrator for assistance.

**Important:** This procedure is disruptive to the Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication services. For Tivoli Storage Productivity Center for Replication, you can perform a takeover to a standby server to ensure that consistency can be maintained for any active servers while the change to security is made.

1. Before changing the user authentication method, stop and restart only the Tivoli Storage Productivity Center Device server before running this procedure.
2. Back up the Device server, Tivoli Integrated Portal, and Tivoli Storage Productivity Center for Replication WebSphere configuration files. The WebSphere configuration backup files provide a restoration point in the event of a failure when performing this procedure.

#### **Back up WebSphere configuration files on UNIX or Linux systems**

Run the following commands:

```
/IBM/Tivoli/tip/profiles/TIPProfile/bin/backupConfig.sh
/IBM/TPC/device/apps/was/profiles/deviceServer/bin/backupConfig.sh
/IBM/replication/eWAS/profiles/CSM/bin/backupConfig.sh
```

#### **Back up WebSphere configuration files on Windows systems**

Run the following commands:

```
\IBM\Tivoli\tip\profiles\TIPProfile\bin\backupConfig.bat
\IBM\TPC\device\apps\was\profiles\deviceServer\bin\backupConfig.bat
\IBM\replication\eWAS\profiles\CSM\bin\backupConfig.bat
```

#### **Back up the soap.client.props file for Tivoli Integrated Portal, Tivoli Storage Productivity Center, and Tivoli Storage Productivity Center for Replication on Windows system**



Use the Windows **copy** command to back up these files:

```
\IBM\Tivoli\tip\profiles\TIPProfile\properties\soap.client.props
\IBM\TPC\device\apps\was\profiles\deviceServer\properties\soap.client.props
\IBM\replication\ewAS\profiles\CSM\properties\soap.client.props
```

### **Back up the soap.client.props file for Tivoli Integrated Portal, Tivoli Storage Productivity Center, and Tivoli Storage Productivity Center for Replication on UNIX or Linux system**

Use the UNIX or Linux **cp** command to back up these files:

```
/IBM/Tivoli/tip/profiles/TIPProfile/properties/soap.client.props
/IBM/TPC/device/apps/was/profiles/deviceServer/properties/soap.client.props
/IBM/replication/ewAS/profiles/CSM/properties/soap.client.props
```

The WebSphere backup command generates a zip file: WebSphereConfig\_YYYY-DD-MM.zip. When you run the backup command, the WebSphere servers will shut down. The WebSphere servers need to be restarted after the backup command is finished. You can restore the WebSphere configuration file by running the **restoreConfig.sh** (for UNIX or Linux) or **restoreConfig.bat** (for Windows) command.

**Note:** If the Tivoli Integrated Portal server fails to start after you run the backup command and the WebSphere servers restart, you must start the Tivoli Integrated Portal server. To start the server, follow these steps:

- For Windows, click **Start > Control Panel > Administrative Tools > Services**. On the Windows Services panel, right-click **Tivoli Integrated Portal - TIPProfile\_Port\_xxxx** where *xxxx* indicates the port specified during installation. Click **Start**.
- For Linux or AIX, go to a command prompt window and enter the following command:

```
TIP_installation_directory/tip/profiles/TIPProfile/bin/
startServer.sh server1
```

*TIP\_installation\_directory* is where Tivoli Integrated Portal is installed.

3. Start an Internet Explorer or Firefox Web browser, and type the following information in the address bar:

```
http://hostname:port
```

Where *hostname* defines the server that is running Tivoli Integrated Portal such as the server name or IP address and *port* defines the port number for Tivoli Integrated Portal. If the default port was accepted during the installation of Tivoli Integrated Portal, the port number is 16310. Contact your Tivoli Storage Productivity Center administrator if you need to verify the host name and port number.

4. On the Tivoli Integrated Portal logon page, log on using the appropriate user ID and password. Your user ID must have administrator permissions.
5. In the Tivoli Integrated Portal navigation tree, click **Security > Secure administration, applications, and infrastructure**.
6. On the Secure administration, applications, and infrastructure page, select **Federated Repositories** from the **Available Realm Definitions** list.

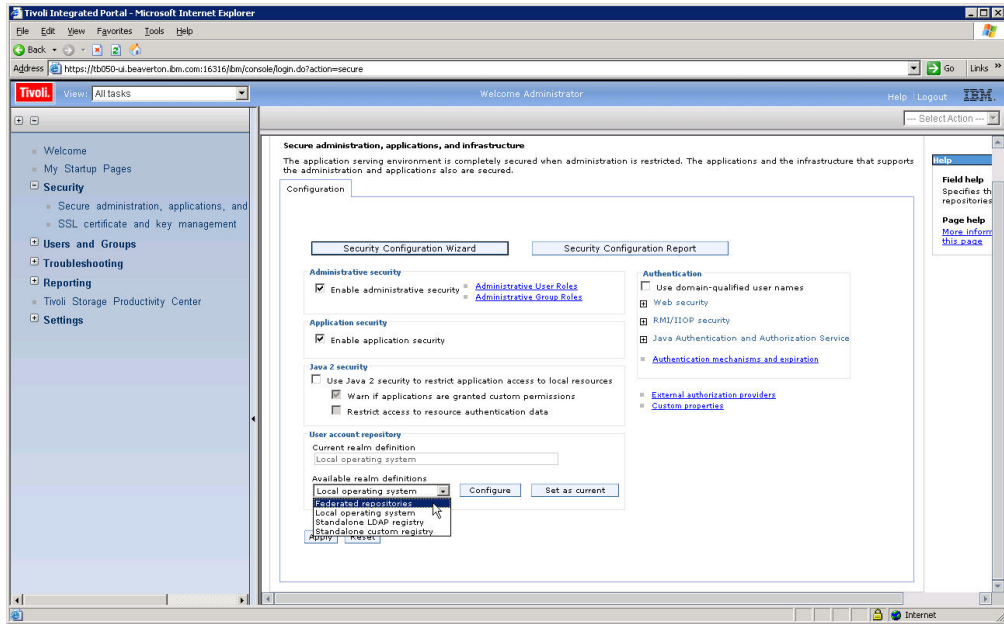


Figure 10. Selecting Federated repositories

7. Click **Configure**. The Federated repositories page is displayed.
8. Under **Related Items**, click **Manage repositories**.

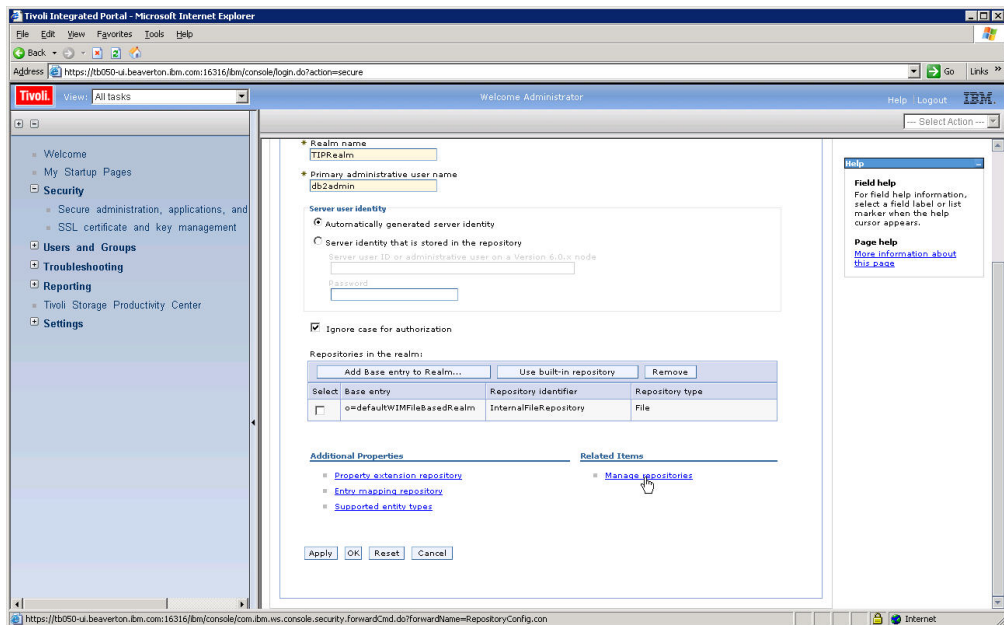


Figure 11. Manage repositories

9. On the Manage repositories page, add the LDAP repository that you want to use for authentication as follows:
  - a. Click **Add** to add a new repository.

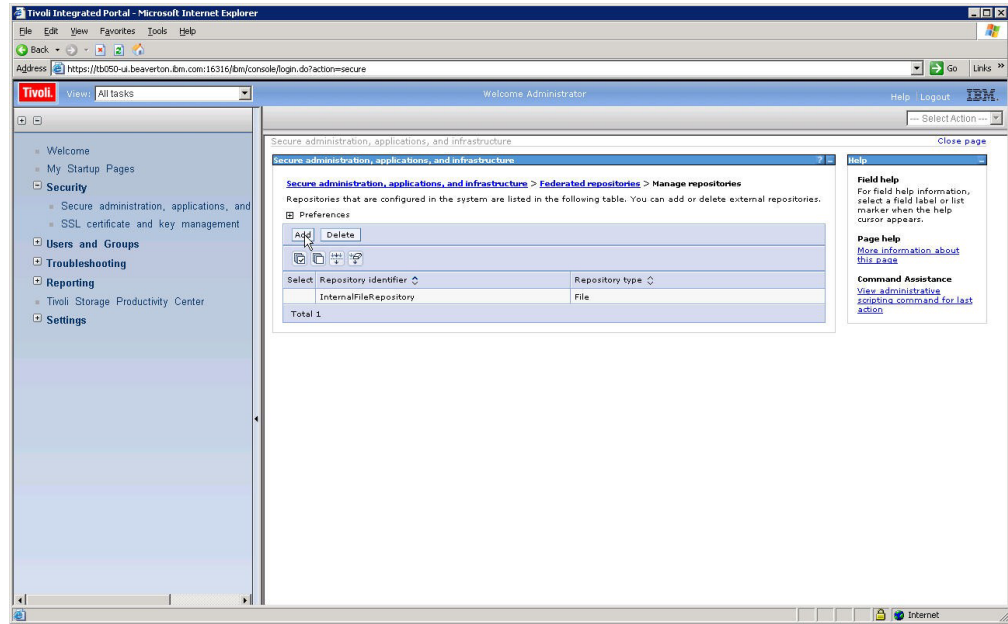


Figure 12. Adding a new repository

- b. Enter the values for the following fields:
  - **Repository identifier.** A unique identifier for the LDAP repository. This identifier uniquely identifies the repository within the cell, for example: LDAP1.
  - **Directory type.** The type of LDAP server to which you want to connect.
  - **Primary host name.** The host name of the primary LDAP server. This host name is either an IP address or a domain name service (DNS) name.
  - **Port.** The LDAP server port. The default value is 389, which is not a Secure Sockets Layer (SSL) connection. You can use port 636 for a Secure Sockets Layer (SSL) connection. For some LDAP servers, you can specify a different port for a non-SSL or SSL connection. If you do not know the port to use, contact your LDAP server administrator.
  - **Bind distinguished name.** The distinguished name (DN) for the application server to use when binding to the LDAP repository. If no name is specified, the application server binds anonymously. In most cases, bind DN and bind password are needed. However, when anonymous bind can satisfy all the required functions, bind DN and bind password are not needed.
  - **Bind password.** The password for the application server to use when binding to the LDAP repository.
- c. Click **OK**.

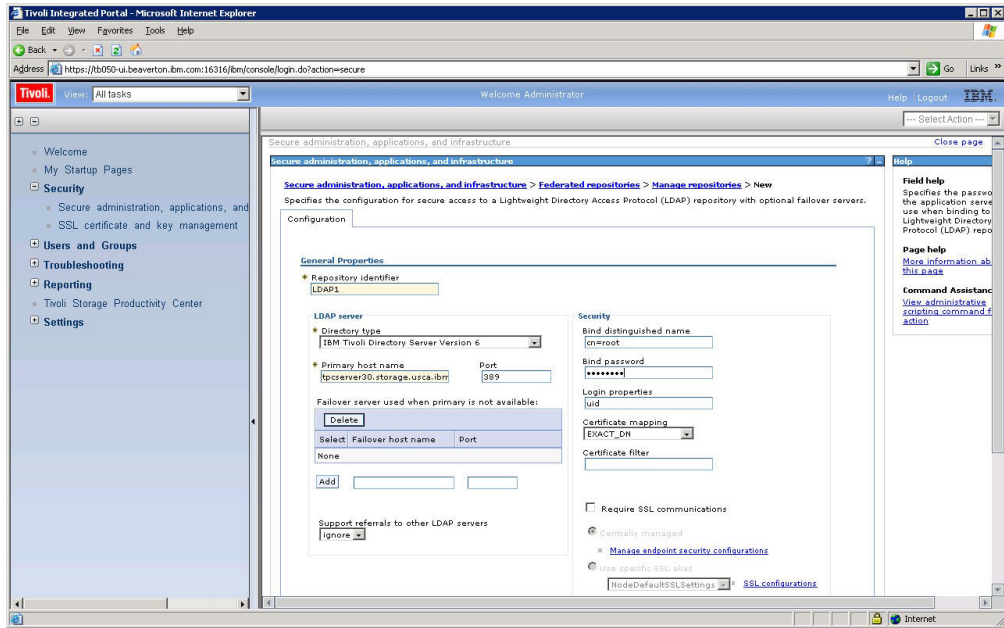


Figure 13. Enter values for a new repository

- d. In the Messages box on the Manage repositories page, click the Save link in Save to the master configuration.
10. From the Manage repositories page, return to the Federated repositories page.
11. Under Repositories in the realm, click Add base entry to Realm.

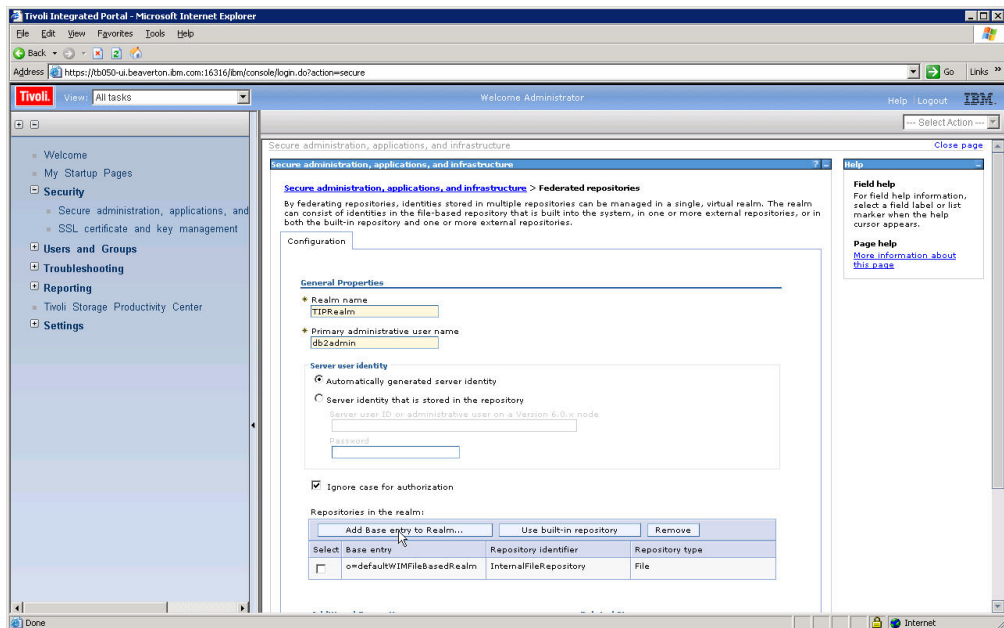


Figure 14. Add base entry to realm

12. On the Repository reference page, configure the following items:
  - a. In the Repository list, select the repository that you created in 9 on page 228.

- b. In the **Distinguished name of a base entry that uniquely identifies this set of entries in the realm** field, enter the distinguished name of a base entry that uniquely identifies the repository in the realm.
- c. In the **Distinguished name of a base entry in this repository** field, enter the distinguished name of the base entry within the repository. In most instances, this value is the same value that you entered in the **Distinguished name of a base entry that uniquely identifies this set of entries in the realm** field.
- d. Click **OK** and then click the **Save** link in **Save to the master configuration**.

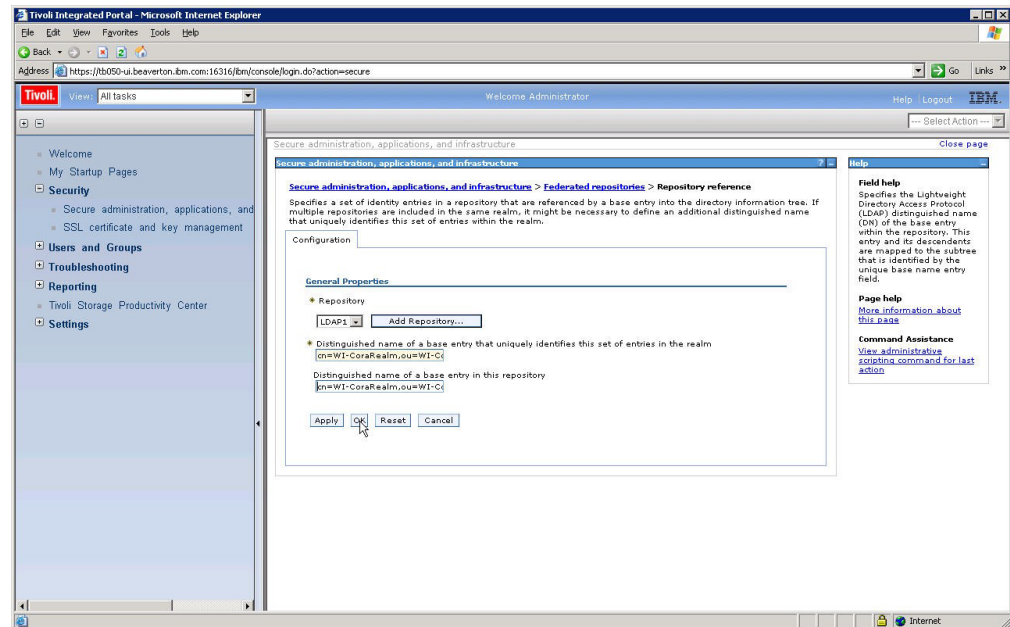


Figure 15. Enter values for a new repository continued

13. On the **Federated repositories** page, there are now two repositories that are displayed under **Repositories in the realm**: the repository that you have added and a default repository that shows **File** in the **Repository type** column. Configure the following items:
  - a. Leave the value in the **Realm name** field as is or change the name of the realm name.
  - b. In the **Primary administrative user name** field, enter the name of a user in the repository that you added. This user is granted administrative privileges in the Tivoli Integrated Portal server, the Tivoli Storage Productivity Center Device server, and the IBM Tivoli Storage Productivity Center for Replication server.
  - c. Click **Server identity that is stored in the repository**.
  - d. In the **Server user ID or administrative user on a Version 6.0.x node** field, enter the same ID that you entered in the **Primary administrative user name** field and enter the password for the user ID in the **Password** field.
  - e. Click **OK** and then click the **Save** link in **Save to the master configuration**.

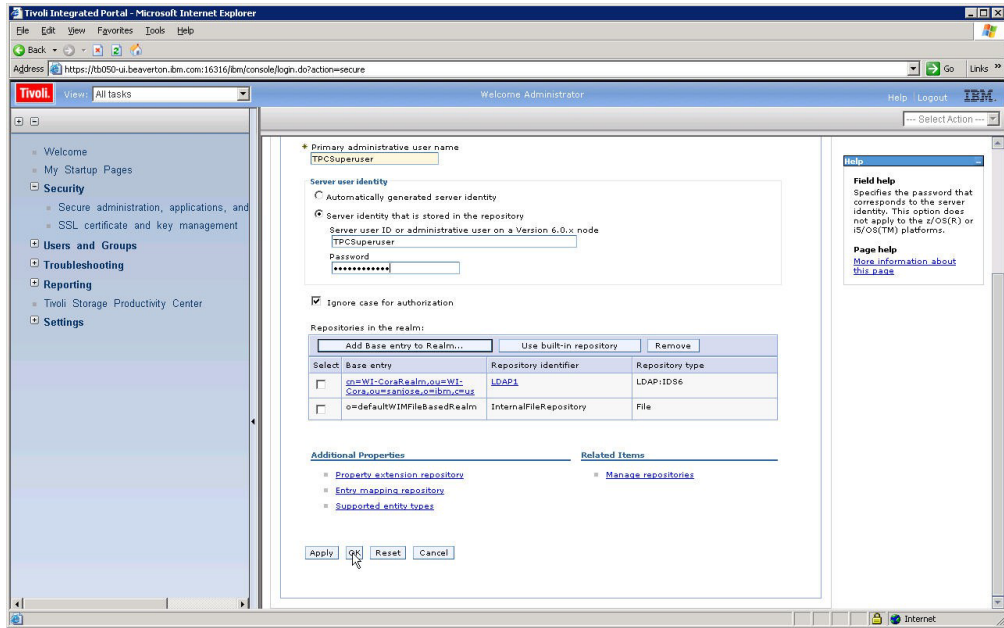


Figure 16. Configuring Federated repository

14. Return to Secure administration, applications, and infrastructure page and configure the following items:
  - a. In the **Available realm definitions** list, select **Federated repositories** and then click **Set as current**.

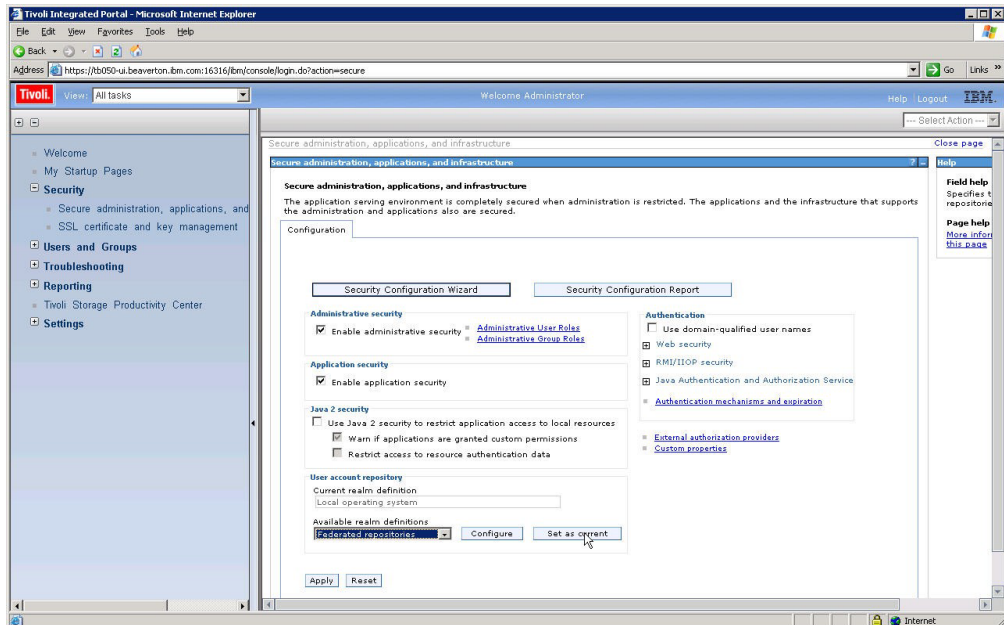


Figure 17. Configuring Federated repository continued

- b. Click **Apply** and then click the **Save** link in **Save to the master configuration**
    - c. Log off Tivoli Integrated Portal.

15. Stop and restart the Tivoli Integrated Portal, IBM Tivoli Storage Productivity Center for Replication servers and the IBM Tivoli Storage Productivity Center Data and Device servers.

**Important:** This step is disruptive to the Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication services. For Tivoli Storage Productivity Center for Replication, you can perform a takeover to a standby server to ensure that consistency can be maintained for any active servers while the change to this procedure is made.

Stop the servers in this order:

- a. Tivoli Storage Productivity Center for Replication server
- b. Tivoli Storage Productivity Center Data server
- c. Tivoli Storage Productivity Center Device server
- d. Tivoli Integrated Portal server

Start the servers in this order:

- a. Tivoli Integrated Portal server
- b. Tivoli Storage Productivity Center Device server
- c. Tivoli Storage Productivity Center Data server
- d. Tivoli Storage Productivity Center for Replication server

For information regarding stopping and starting a server, go to the Tivoli Storage Productivity Center Information Center and search for *starting and stopping services*. If the server is running on a Windows operating system, the change to the user authentication method might not be implemented until you restart the computer.

To verify that the federated repository is configured correctly, complete the following steps:

1. log on to Tivoli Integrated Portal as the new super user ID (the administrative account is no longer the OS credentials, but is now configured to use the LDAP server super user credentials).
2. In the Tivoli Integrated Portal navigation tree, click **Users and Groups > Manage Users** .
3. In the **Search by** list, select **User ID**.
4. Click **Search** to search for users in the federated repository. The list of users includes users from both the LDAP and the local file registry.
5. After you have completed and verified the conversion of your Tivoli Storage Productivity Center system's authentication mechanism from local operating system to LDAP federated repositories, remove all Administrative Group Roles entries in the Tivoli Integrated Portal console that involve local operating system groups:
  - a. In the Tivoli Integrated Portal navigation tree, click **Users and Groups > Administrative Group Roles**.
  - b. In the Administrative Group Roles table, select all entries that involve local operating system groups. Do **not** select entries that involve LDAP federated repository groups.
  - c. Click **Remove**.
  - d. In the Messages box on the Administrative Group Roles page, click **Save link** in "Save directly to the master configuration."

After you change the authentication method successfully, you can assign roles to individuals from the Role-to-Group Mapping node in Tivoli Storage Productivity

Center graphical user interface (GUI). You can map Tivoli Storage Productivity Center roles, such as Tape Operator or Fabric Administrator, to user groups that you create in your LDAP-compliant repository. To learn more, go to the IBM Tivoli Storage Productivity Center Information Center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp> and search for *Role-to-Group Mappings*.

**Related tasks:**

“Enabling secure communication between Tivoli Storage Productivity Center and the LDAP repository”

You can use the Secure Socket Layer (SSL) protocol to secure the communication between IBM Tivoli Storage Productivity Center and the LDAP repository that you are using for user authentication. The SSL protocol provides security and data integrity for communications over Transmission Control Protocol/Internet Protocol (TCP/IP) networks.

“Disabling secure communication between Tivoli Storage Productivity Center and the LDAP repository” on page 236

You can disable the Secure Socket Layer (SSL) protocol between the LDAP repository and the IBM Tivoli Storage Productivity Center system at any time.

**Enabling secure communication between Tivoli Storage Productivity Center and the LDAP repository:**

You can use the Secure Socket Layer (SSL) protocol to secure the communication between IBM Tivoli Storage Productivity Center and the LDAP repository that you are using for user authentication. The SSL protocol provides security and data integrity for communications over Transmission Control Protocol/Internet Protocol (TCP/IP) networks.

To configure for SSL, you must complete the following steps in IBM Tivoli Integrated Portal:

1. Start an Internet Explorer or Firefox web browser, and type the following information in the address bar:

`http://hostname:port`

Where *hostname* defines the server that is running Tivoli Integrated Portal such as the server name or IP address and *port* defines the port number for Tivoli Integrated Portal. If the default port was accepted during the installation of Tivoli Integrated Portal, the port number is 16310. Contact your Tivoli Storage Productivity Center administrator if you must verify the host name and port number.

2. On the Tivoli Integrated Portal logon page, log on using the appropriate user ID and password. Your user ID must have administrator permissions.
3. In the Tivoli Integrated Portal navigation tree, click **Security > SSL certificate and key management > Key stores and certificates > NodeDefaultTrustStore > Signer certificates**.
4. On Signer certificates page, click **Retrieve from port**.
5. On the Retrieve from port page, enter values for the following fields:
  - **Host**. The fully qualified host and domain name of your LDAP-compliant repository.
  - **Port**. The port where your LDAP repository is listening for secure communications; this is usually port 636.
  - **SSL configuration for outbound connection**. Accept the default value.
  - **Alias**. An alias name for the retrieved certificate.
6. Click **Retrieve signer information**.



7. When the signer information is displayed, click **OK**.
8. In the **Messages** box on the Signer certificates page, click the **Save** link in **Save to the master configuration**.
9. On the Signer certificates page, select the LDAP repository certificate that you retrieved and click **Extract**.
10. On the Extract signer certificate page, enter values for the following fields:
  - **File name.** The file name for this certificate. For example, LDAPSSLCert. This file is saved in the C:\Tivoli Integrated Portal install directory\profiles\TIPProfile\etc\ directory on the Windows platform. In the AIX, Linux, and UNIX platforms, this file is saved in the /Tivoli Integrated Portal install directory/profiles/TIPProfile/etc/ directory. Tivoli Integrated Portal install directory is the directory where Tivoli Integrated Portal is installed.
  - **Data type.** Accept the default value.
11. Click **OK**. A message is displayed on the Signer certificates page stating that the file was successfully extracted.
12. In the Tivoli Integrated Portal navigation tree, click **Security > Secure administration, applications, and infrastructure**.
13. On the Secure administration, applications, and infrastructure page, make sure that **Federated Repositories** is selected in the **Available Realm Definitions** list.
14. Click **Configure**. The Federated repositories page is displayed.
15. Under **Related Items**, click **Manage repositories**.
16. On the Manage repositories page, click the identifier for the repository for which you want to enable the SSL protocol.
17. On the configuration page for the repository, configure the following items:
  - a. In the **Port** field, enter 636
  - b. Select the **Require SSL communications** check box.
18. Click **OK**.
19. In the **Messages** box on the Manage repositories page, click the **Save** link in **Save to the master configuration**.
20. Log off from Tivoli Integrated Portal.
21. From the command line, stop the Tivoli Integrated Portal server, the Tivoli Storage Productivity Center for Replication, and the Tivoli Storage Productivity Center Device and Data servers. For information regarding starting and stopping a server, go to the Tivoli Storage Productivity Center Information Center and search for *starting and stopping services*.
22. Use the following commands to import the LDAP repository certificate into the Tivoli Storage Productivity Center Device server keystore:

#### **Windows**

```
cd C:\Tivoli Storage Productivity Center install directory\jre\bin\
keytool -import -keystore C:\Tivoli Storage Productivity Center install directory
\device\apps\was\profiles\deviceServer\config\cells\DefaultNode\nodes\DefaultNode\
trust.p12 -storetype pkcs12 -storepass WebAS -file
C:\Tivoli Integrated Portal install directory\profiles\TIPProfile\etc\
LDAP repository certificate file name
```

where *Tivoli Storage Productivity Center install directory* is the directory where Tivoli Storage Productivity Center is installed, *Tivoli Integrated Portal install directory* is the directory where Tivoli Integrated Portal is installed, and *LDAP repository certificate file name* is the file name that you assigned to the LDAP repository certificate in Step 11.

Type yes when you are asked if you want to trust this certificate.

### AIX, Linux, and UNIX

```
cd Tivoli Storage Productivity Center install directory/jre/bin/
./keytool -import -keystore /Tivoli Storage Productivity Center
install directory/device/apps/was/profiles/deviceServer/config/
cells/DefaultNode/nodes/DefaultNode/trust.p12 -storetype pkcs12
-storepass WebAS -file /Tivoli Storage Productivity Center
install directory/profiles/TIPProfile/etc/
LDAP repository certificate file name
```

where *Tivoli Storage Productivity Center install directory* is the directory where Tivoli Storage Productivity Center is installed, *Tivoli Integrated Portal install directory* is the directory where Tivoli Integrated Portal is installed, and *LDAP repository certificate file name* is the file name that you assigned to the LDAP repository certificate in Step 11.

Type yes when you are asked if you want to trust this certificate.

23. Use the following commands to import the LDAP repository certificate into the Tivoli Storage Productivity Center for Replication server keystore:

### Windows

```
keytool -import -keystore <C:Tivoli Storage Productivity Center install
directory>\eWAS\profiles\CSM\config\cel\ls\DefaultNode\nodes\DefaultNode
\trust.p12 -storetype pkcs12 -storepass WebAS -file <C:Tivoli
Storage Productivity Center install directory>\profiles\TIPProfile\
etc\LDAP repository certificate file name
```

where *Tivoli Storage Productivity Center for Replication install directory* is the directory where Tivoli Storage Productivity Center for Replication is installed.

Type yes when you are asked if you want to trust this certificate.

### AIX, Linux, and UNIX

```
./keytool -import -keystore /Tivoli Storage Productivity Center install
directory/eWAS/profiles/CSM/config/cells/DefaultNode/nodes/DefaultNode
/trust.p12 -storetype pkcs12 -storepass WebAS -file /Tivoli Storage
Productivity Center install directory/profiles/TIPProfile/etc/
LDAP repository certificate file name
```

where *Tivoli Storage Productivity Center for Replication install directory* is the directory where Tivoli Storage Productivity Center for Replication is installed.

Type yes when you are asked if you want to trust this certificate.

24. From the command line, start the Tivoli Integrated Portal server, the Tivoli Storage Productivity Center for Replication, and the Tivoli Storage Productivity Center Device and Data servers.

**Important:** Securing the communication between the LDAP server and the Tivoli Storage Productivity Center for Replication system modifies the LDAP settings. When the LDAP settings are modified, all of the group-to-role mappings in Tivoli Storage Productivity Center for Replication are reset to the Tivoli Integrated Portal administrator ID.

You must manually re-create the Tivoli Storage Productivity Center for Replication group-to-role mappings. The Tivoli Storage Productivity Center and Tivoli Integrated Portal group-to-role mappings are preserved and require no adjustment.

### Disabling secure communication between Tivoli Storage Productivity Center and the LDAP repository:

You can disable the Secure Socket Layer (SSL) protocol between the LDAP repository and the IBM Tivoli Storage Productivity Center system at any time.

To disable the use of SSL, you must complete the following steps in IBM Tivoli Integrated Portal:

1. Start an Internet Explorer or Firefox web browser, and type the following information in the address bar:

`http://hostname:port`

Where *hostname* defines the server that is running Tivoli Integrated Portal such as the server name or IP address and *port* defines the port number for Tivoli Integrated Portal. If the default port was accepted during the installation of Tivoli Integrated Portal, the port number is 16310. Contact your Tivoli Storage Productivity Center administrator if you must verify the host name and port number.

2. On the Tivoli Integrated Portal logon page, log on using the appropriate user ID and password. Your user ID must have administrator permissions.
3. In the Tivoli Integrated Portal navigation tree, click **Security > Secure administration, applications, and infrastructure**.
4. On the Secure administration, applications, and infrastructure page, make sure that **Federated Repositories** is selected in the **Available Realm Definitions** list.
5. Click **Configure**. The Federated repositories page is displayed.
6. Under **Related Items**, click **Manage repositories**.
7. On the Manage repositories page, click the identifier for the repository for which you want to disable the SSL protocol.
8. On the configuration page for the repository, configure the following items:
  - a. In the **Port** field, enter 389.
  - b. Clear the **Require SSL communications** check box.
9. Click **OK**.
10. In the **Messages** box on the Manage repositories page, click the **Save** link in **Save to the master configuration**.
11. Log off from Tivoli Integrated Portal.
12. From the command line, stop and then start the Tivoli Integrated Portal server, the Tivoli Storage Productivity Center for Replication, and the Tivoli Storage Productivity Center Device and Data servers. For information regarding starting and stopping a server, go to the Tivoli Storage Productivity Center Information Center and search for *starting and stopping services*.

**Important:** Disabling the communication between the LDAP server and the Tivoli Storage Productivity Center for Replication system modifies the LDAP settings. When the LDAP settings are modified, all of the group-to-role mappings in Tivoli Storage Productivity Center for Replication are reset to the Tivoli Integrated Portal administrator ID.

You must manually re-create the Tivoli Storage Productivity Center for Replication group-to-role mappings. The Tivoli Storage Productivity Center and Tivoli Integrated Portal group-to-role mappings are preserved and require no adjustment.

### **Changing the authentication method from LDAP to local operating system**

You can configure IBM Tivoli Storage Productivity Center to authenticate against the local operating system. To change the authentication method from LDAP to local operating system, you must use Tivoli Integrated Portal.

If you use operating system authentication, the use of the single sign-on feature is limited. Operating system authentication does not support single sign-on when Tivoli Storage Productivity Center and Tivoli Integrated Portal are on separate computers. Operating system authentication also does not support single sign-on for element managers regardless of the location of the element manager application.

**Important:** This procedure is disruptive to the Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication services. For Tivoli Storage Productivity Center for Replication, you can complete a takeover to a standby server to ensure that consistency can be maintained for any active servers while the change to security is finished.

To change the user authentication method from LDAP to local operating system, complete the following steps:

1. Before you change the user authentication method, stop and restart only the Tivoli Storage Productivity Center Device server before you complete this procedure.
2. Back up the Device server and Tivoli Integrated Portal WebSphere configuration files. The WebSphere configuration backup files provide a restoration point in the event of a failure when you complete this procedure.

#### **Back up WebSphere configuration files on Windows systems**

Run the following commands:

```
\IBM\Tivoli\tip\profiles\TIPProfile\bin\backupConfig.bat
\IBM\TPC\device\apps\was\profiles\deviceServer\bin\backupConfig.bat
\IBM\replication\eWAS\profiles\CSM\bin\backupConfig.bat
```

#### **Back up WebSphere configuration files on UNIX or Linux systems**

Run the following commands:

```
/IBM/Tivoli/tip/profiles/TIPProfile/bin/backupConfig.sh
/IBM/TPC/device/apps/was/profiles/deviceServer/bin/backupConfig.sh
/IBM/replication/eWAS/profiles/CSM/bin/backupConfig.sh
```

#### **Back up the soap.client.props file for Tivoli Integrated Portal, Tivoli Storage Productivity Center, and Tivoli Storage Productivity Center for Replication on Windows system**

Use the Windows **copy** command to back up these files:

```
\IBM\Tivoli\tip\profiles\TIPProfile\properties\soap.client.props
\IBM\TPC\device\apps\was\profiles\deviceServer\properties\soap.client.props
\IBM\replication\eWAS\profiles\CSM\properties\soap.client.props
```

#### **Back up the soap.client.props file for Tivoli Integrated Portal, Tivoli Storage Productivity Center, and Tivoli Storage Productivity Center for Replication on UNIX or Linux system**

Use the UNIX or Linux **cp** command to back up these files:

```
/IBM/Tivoli/tip/profiles/TIPProfile/properties/soap.client.props
/IBM/TPC/device/apps/was/profiles/deviceServer/properties/soap.client.props
/IBM/replication/eWAS/profiles/CSM/properties/soap.client.props
```

The WebSphere backup command generates a compressed file:

WebSphereConfig\_YYYY-DD-MM.zip. When you run the backup command, the WebSphere servers shut down. The WebSphere servers must be restarted after the backup command is finished. You can restore the WebSphere configuration file by running the **restoreConfig.sh** (for UNIX or Linux) or **restoreConfig.bat** (for Windows) command.

3. Start an Internet Explorer or Firefox web browser, and type the following information in the address bar:

http://hostname:port

Where *hostname* defines the server that is running Tivoli Integrated Portal such as the server name or IP address and *port* defines the port number for Tivoli Integrated Portal. If the default port was accepted during the installation of Tivoli Integrated Portal, the port number is 16310. Contact your Tivoli Storage Productivity Center administrator to verify the host name and port number.

4. On the Tivoli Integrated Portal logon page, log on using the appropriate user ID and password. Your user ID must have administrator permissions.
5. In the Tivoli Integrated Portal navigation tree, click **Security > Secure administration, applications, and infrastructure**.
6. On the Secure administration, applications, and infrastructure page, select **Local operating system** from the **Available Realm Definitions** list.

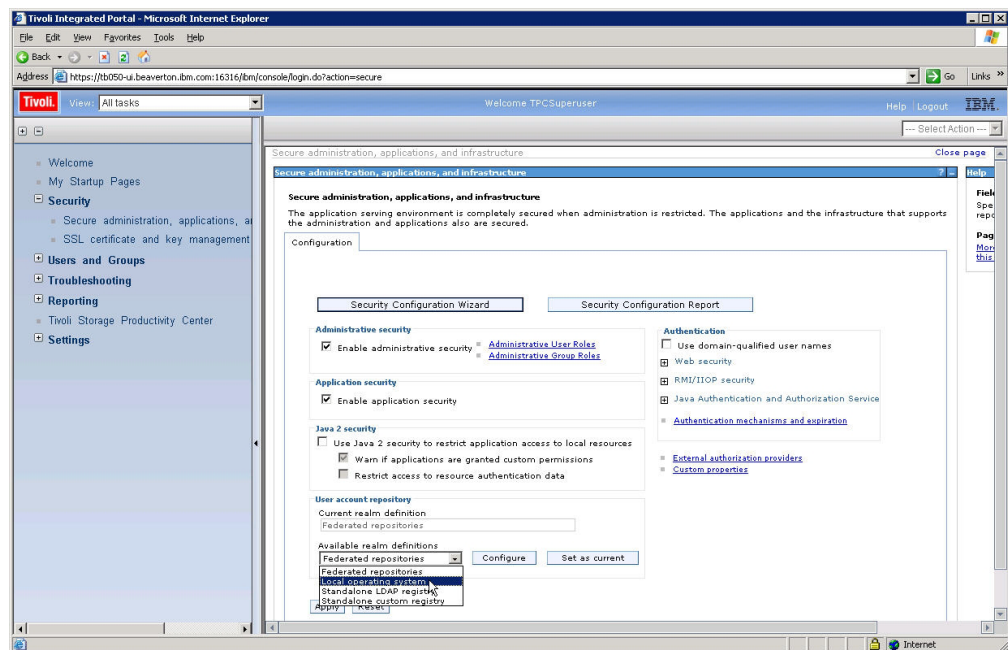


Figure 18. Selecting local operating system

7. Click **Configure**. The Local operating system page is displayed.
8. On the Local operating system page, configure the following items:
  - a. In the **Primary administrative user name** field, enter the name of a user that is defined in your local operating system. This user is granted administrative privileges in the following components:
    - Tivoli Integrated Portal server
    - Tivoli Storage Productivity Center Device server
    - IBM Tivoli Storage Productivity Center for Replication server

**Note:** When you specify a name in the **Primary administrative user name** field, do not specify the domain. If you specify the domain, when IBM WebSphere Application Server instances (such as Tivoli Integrated Portal or Tivoli Storage Productivity Center for Replication) are synchronized with Tivoli Integrated Portal authentication configuration and may not shut down correctly.

- b. Click **Server identity** that is stored in the repository.
- c. In the **Server user ID or administrative user on a Version 6.0.x node** field, enter the same ID that you entered in the **Primary administrative user name** field and enter the password for the user ID in the **Password** field.
- d. Click **OK**.

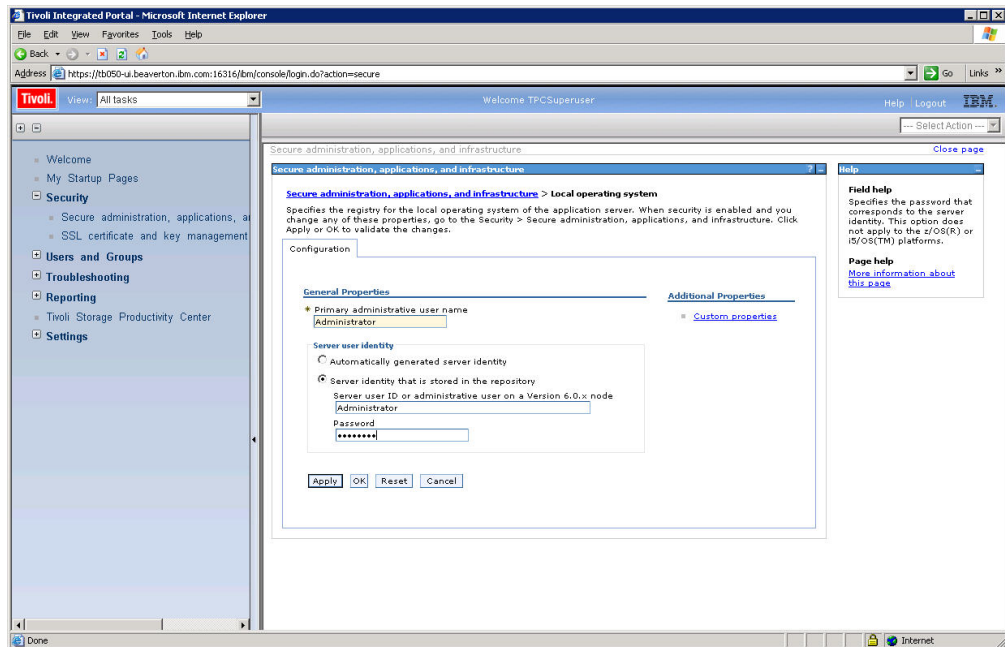


Figure 19. Configuring local operating system

- e. In the **Messages** box on the Secure administration, applications, and infrastructure page, click the **Save** link in **Save to the master configuration**.
9. Configure the following items on the Secure administration, applications, and infrastructure page:
    - a. In the **Available realm definitions** list, select **Local operating system** and then click **Set as current**.

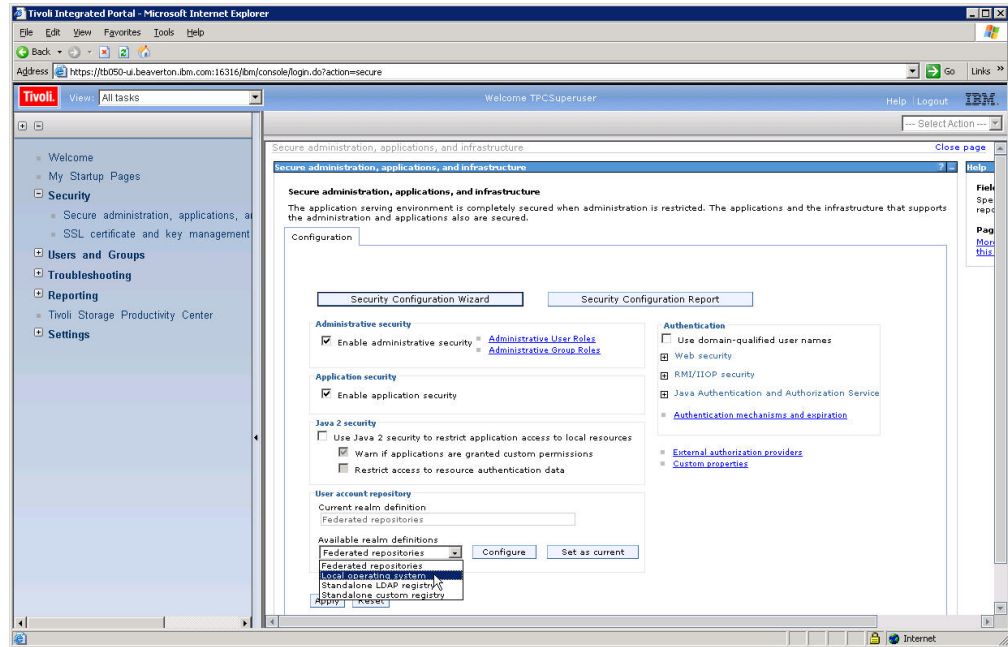


Figure 20. Selecting local operating system for available realm definition

- b. Click **Apply** and then click the **Save** link in **Save to the master configuration**
  - c. Log out from IBM Tivoli Integrated Portal.
10. Stop and restart the Tivoli Integrated Portal, IBM Tivoli Storage Productivity Center for Replication servers and the IBM Tivoli Storage Productivity Center Data and Device servers.

**Important:** This step is disruptive to the Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication services. For Tivoli Storage Productivity Center for Replication, you can complete a takeover to a standby server and ensure that consistency is maintained for any active servers while the change to this procedure has finished.

Stop the servers in this order:

- a. Tivoli Storage Productivity Center for Replication server
- b. Tivoli Storage Productivity Center Data server
- c. Tivoli Storage Productivity Center Device server
- d. Tivoli Integrated Portal server

Start the servers in this order:

- a. Tivoli Integrated Portal server
- b. Tivoli Storage Productivity Center Device server
- c. Tivoli Storage Productivity Center Data server
- d. Tivoli Storage Productivity Center for Replication server

For information about stopping and starting a server, go to the Tivoli Storage Productivity Center Information Center and search for *starting and stopping services*. If the server is running on a Windows operating system, the change to the user authentication method might not be implemented until you restart the computer. A valid group from LDAP must be mapped to the Tivoli Storage Productivity Center Superuser role. For information about mapping a

user role in Tivoli Storage Productivity Center, see <http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp> and search for **Role-to-Group Mappings**.

## Changing the LTPA token expiration for single sign-on

For single sign-on, the user's authentication information is passed between applications using Lightweight Third-Party Authentication (LTPA) tokens. By default, the tokens expire after 24 hours. You can change the LTPA token expiration time using the Tivoli Integrated Portal user interface.

To change the expiration time for the LTPA tokens, complete the following steps.

1. Start an Internet Explorer or Firefox Web browser, and type the following information in the address bar:

`http://hostname:port`

where *hostname* defines the server that is running Tivoli Integrated Portal such as server name or IP address and *port* defines the port number for Tivoli Integrated Portal. If you accepted the default port during the installation of Tivoli Integrated Portal, the port number is 16310. Contact your Tivoli Storage Productivity Center administrator if you need to verify the host name and port number.

2. On the Tivoli Integrated Portal logon page, log on using the appropriate user ID and password. Your user ID must have administrator permissions.
3. In the Tivoli Integrated Portal navigation tree, click **Security > Secure administration, applications, and infrastructure**.
4. Under **Authentication**, click **Authentication mechanisms and expiration**.
5. On the **Configuration** tab under **Authentication expiration**, type the expiration time in minutes in the **Timeout value for forwarded credentials between servers** field. The expiration time must be greater than the authentication cache timeout value that is shown above the **Timeout value for forwarded credentials between servers** field. The default expiration time is 1440 minutes (24 hours).
6. Click **Apply**.
7. In the Messages box, click **Save directly to the master configuration**.

---

## Administering data sources

Use the following topics for information about how to administer CIMOM agents, Data agents, Storage Resource agents, Inband Fabric agents, Out of Band Fabric agents, Storage Subsystems, TPC Servers, and VMware VI data sources.

### Checking the health of your data sources

This topic provides information about checking the health of your data sources.

The storage systems that use the native interfaces (for example, DS8000, the XIV system, SAN Volume Controller, and Storwize V7000) do not use CIMOM agents. Once you successfully connect to the storage system, Tivoli Storage Productivity Center can display the status.

Expand **Administrative Services > Data Sources** and left-click the following nodes:

- **CIMOM Agents**
- **Data/Storage Resource Agents**



- **Inband Fabric Agents**
- **Out of Band Fabric Agents**
- **Storage Subsystems**
- **TPC Servers**
- **VMware VI Data Source**

A list of data sources is displayed in the content pane. If the agents are up and running, there is a green health status icon in the State column. For the CIMOM Agents and Storage Subsystems, a green health status icon is displayed in the Connection Status column.

## Manually changing the Windows service logon for the agent

This topic describes how to change the Windows service logon for the Common agent.

To change the Windows service logon for the Common agent, complete the following steps:

1. In Windows, open the Services panel: **Start > Settings > Control Panel > Administrative Tools > Services**.
2. On the Services panel, right-click **IBM Tivoli Common agent - 'C:\Program Files\IBM\TPC\ca'**. Select **Properties**.
3. Select the Log On tab. Change **This account** field with your changed logon ID. If your Tivoli Storage Productivity Center server is part of a Windows domain, change this logon to <domain>\<account>. For example, **mydomain\myaccount**. Enter the password if you have changed the password. Click **Enable** and then **OK**. The Common agent requires that the domain account have local administrator privileges as well as "Log on as a service" and "Act as part of the operating system" user rights.

## CIM agents

Common Information Model (CIM) agents provide a CIM interface for storage entities. IBM Tivoli Storage Productivity Center uses CIM agents to collect information about the entities that you want to monitor.

These entities include IBM TotalStorage Enterprise Storage Server, IBM System Storage DS4000, IBM System Storage DS5000 Storage Manager, IBM System Storage DS6000, tape, switch, and any SMI-S certified Common Information Model Object Manager (CIMOM).

CIM agents are no longer required for DS8000, XIV, and SAN Volume Controller storage systems.

Before you can monitor the entities that are managed by CIM agents, you must configure the agents in the Tivoli Storage Productivity Center user interface. Use either of the following methods to add a CIM agent:

- Use the Configure Devices wizard to automatically add a CIM agent and discover the entities that it manages. See "Configure Devices" on page 305 for more information.
- Manually add the CIM agent using the **CIMOM Agents** node under **Administrative Services > Data Sources** in the navigation tree. See "Adding a CIM agent connection" on page 244 for more information.

After you manually add a CIM agent, you must run a discovery to detect the entities that it manages. See "Discover storage subsystem, tape library, and fabric

information with CIM agents” on page 330 and “Discover CIM agents automatically” on page 245 for more information about the discovery process.

## Adding a CIM agent connection

CIM agents provide a Common Information Model (CIM) interface for management applications. This topic discusses how to add a CIM agent connection.

You can add a CIM agent connection to Tivoli Storage Productivity Center by using the Configure Devices wizard or use the manual steps in this procedure.

To add a CIM agent connection using the Configure Devices wizard, see “Configure Devices” on page 305.

To add a CIM agent connection to Tivoli Storage Productivity Center manually, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services > Data Sources** and click **CIMOM Agents**.
2. In the content pane, click **Add CIMOM**. The Add CIMOM window opens.
3. Specify information about the CIM agent:

**Host** The hostname of the computer on which the CIMOM is installed. For element managers, you can enter a host name, IPv4, or IPv6 address depending on what is supported within your environment:

- Enter an IPv4 address in the following format: `###.###.###.###`. For example: 192.168.1.65
- Enter an IPv6 address. The preferred IPv6 address representation is written as eight groups of four hexadecimal digits: `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx`, where each x is a hexadecimal digit representing 4 bits. For example: 2001:DB8:0000:1234:0000:0000:5678:ABCD. You can also specify IPv6 addresses using shortened formats that omit leading zeros: 2001:DB8:0:1234:0:0:5678:ABCD or use double colons: 2001:DB8:0000:1234::5678:ABCD in place of a series of zeros.

You can enter an IPv6 address for a CIMOM if the computer where IBM Tivoli Storage Productivity Center is installed is IPv6 or dual stack (IPv4 and IPv6) enabled.

**Port** The port on which the CIMOM is listening. By default this port is 5989 for a secure connection and 5988 for an unsecured connection. See the documentation for the appropriate SMI-S CIM Provider for more information about available ports.

### Username

The user ID used for authentication, if required by the CIMOM.

### Password

The password, if required by the CIMOM.

### Interoperability Namespace

This namespace within the CIMOM allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how IBM Tivoli Storage Productivity Center interacts with the CIMOM when retrieving information. The following are the default namespaces for CIMOM agents for switches:

- IBM: /root/ibm
- Brocade: /interop

- NetApp: /interop
- Cisco: /root/cimv2  
For version 3.2.1 or later: /root/pg\_interop
- Engenio: /interop
- LSI (Pegasus-based): /root/PG\_InterOp
- EMC: /root/emc
- HDS: For HiCommand 5.6 or higher, use /root/smis/current  
For a HiCommand version that supports SMI-S 1.2, use following namespace to traverse the model as SMI-S 1.1: /root/smis/smis11  
For HiCommand versions lower than HiCommand 5.6, use: /root/hitachi/dmxx, where xx is represents the level of HiCommand.
- HP: /root
- LSI: /root/PG\_InterOp
- McData: /interop
- SUN (storage subsystem): /root/sun3510 or /interop
- XYRATEX: /root/PG\_interop

To view a list of default namespaces for CIMOM agents for storage subsystems that is included with Tivoli Storage Productivity Center, open the following file in a text editor: *install\_directory/data/config/namespace.config*, where *install\_directory* represents the directory where the product is installed.

Check the documentation of the appropriate SMI-S CIM Provider or contact the storage system vendor to ensure that you use the most current namespaces.

**Protocol**

The version of the cim-xml protocol. Can be http or https.

**Display Name**

The name of the CIMOM, as specified by the CIMOM provider, to be displayed in the Tivoli Storage Productivity Center interface.

**Description**

An optional description of the CIMOM. This description can be any combination of alphanumeric characters.

**Test CIMOM connectivity before adding**

Select this option to have Tivoli Storage Productivity Center ensure that communication to the CIMOM is working properly before adding information about that CIMOM.

4. Click **Save**.

This procedure does not physically add a CIMOM, but adds the information you have given about the CIMOM to Tivoli Storage Productivity Center so that Tivoli Storage Productivity Center can communicate with the CIMOM.

**Discover CIM agents automatically**

This topic discusses how to discover storage subsystem CIMOMs that are visible to the Device server using Service Location Protocol (SLP).

In many cases CIM agents can be automatically discovered using SLP. The criteria is that the CIM agent must be visible to the Device server.

To automatically discover all CIM agents that are visible to the Device server, complete the following steps:

1. In the navigation tree pane, expand **Administrative Services > Discovery** and click **CIMOM**.
  - a. Click the **Options** tab.
  - b. Enter the IP address or host name for the SLP directory agents for CIMOM discovery. Or select the **Scan local subnet** check box.
  - c. Click **File > Save** on the menu bar to save the information.
2. Right-click **CIMOM** and click **Run Now**.
3. After you submit a CIMOM discovery job, a message window is displayed with the text **CIMOM Discovery Job submitted**. Click **OK**. The job name is located below **CIMOM**. While it is running, the CIMOM job has a blue circle outline to the left of the job name.
4. To determine if the job has completed, right-click **CIMOM** and click **Update Job Status**. If the job has completed successfully, a green square is shown in front of the job name. If the job has completed but failed, a red circle is shown.

**Note:** The job returns as failed if one or more sub-jobs failed. Some CIMOMs might still have been discovered.

5. After the job completes, the discovered CIM agents are listed in the **CIMOM Agents** table. Expand **Administrative Services > Data Sources** and click **CIMOM Agents**.
6. For each CIM agent that you want to use, highlight the CIM agent and click the magnifying glass at the left of the listing. You must select the CIM agents one at a time. The **CIMOM Agents** window is opened.
7. Enter additional information for the CIM agent, such as the user name and password, and select the **Test CIMOM connectivity before updating** check box.
8. Click **File > Save** on the menu bar to save the information and perform a connectivity check.

**Note:** This procedure does not physically add the CIM agent to your environment, but provides IBM Tivoli Storage Productivity Center with the information it needs to communicate with that CIM agent. This includes information that you have defined and CIM agents discovered using SLP.

## Viewing information about a CIM agent

Use panels and reports in IBM Tivoli Storage Productivity Center to view detailed information about a CIM agent.

Use the following reports to view information about a CIM agent:

- In the Navigation Tree pane, expand **IBM Tivoli Storage Productivity Center > Reporting > Data Source Reports > CIMOM Agents** and click **By CIMOM Agent** or **By Managed Device**. See “Data source reports - CIMOM agents” on page 1093 for more information about these reports.
- In the Navigation Tree pane, expand **Administrative Services > Data Sources** and click **CIMOM Agents**. A list of CIM agents appears in the content pane. Click the magnifying glass next to a CIM agent to view the following information about that agent:

### Service URL

The service URL of the CIMOM containing the IP address of the CIMOM, the port on which the CIMOM is listening, and the protocol

used for communication. This URL has a protocol [http | https], an IP or Hostname, and a port number. This field displays IPv4 and IPv6 addresses as appropriate.

**Display Name**

The name of the CIMOM as specified by the CIMOM provider that will appear in the IBM Tivoli Storage Productivity Center interface.

**Description**

The optional description that was entered on the Add CIMOM window.

**Username**

The CIMOM user ID used for authentication.

**Password**

The password for the CIMOM.

**Interoperability Namespace**

Enter the interoperability namespace of the CIMOM. This namespace within the CIMOM allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how Tivoli Storage Productivity Center interacts with the CIMOM when retrieving information.

**Truststore Location**

The location (path on this computer) of a certificate file for certificate based authentication in the https protocol.

**User Interface Description**

The name of the Human Interface Service (if any) supported by this CIMOM.

**Software Level**

The software version level of the CIMOM agent.

**Protocol Version**

The version of the cim-xml protocol.

**Authentication Mechanism**

The authentication mechanism supported by the CIMOM. This field can contain the following values: Unknown, None, Other, Basic, Digest.

**Alias** The alias of the CIMOM.

**Service ID**

The service ID for the CIMOM.

**Protocol**

The communication protocol used for the CIMOM. Possible values are http and https.

**SLP Attributes**

The standard set of attributes for this CIMOM. The attributes are retrieved via SLP.

**Connection Status**

The status of this CIMOM with respect to Tivoli Storage Productivity Center. Possible values are: SUCCESS, UNCONFIGURED, UNKNOWN, INVALID\_NAMESPACE, TIMEOUT, REFUSED, LOGIN\_FAILED, SSL\_HANDSHAKE\_ERROR, SSL\_REGISTRATION\_INVALID, CIMCLIENT\_ERROR

**Status Timestamp**

The date/time when the Connection Status information was last collected.

**Test CIMOM connectivity before updating**

If this box is checked Tivoli Storage Productivity Center ensures that communication to the CIMOM is working properly before making any changes.

**Updating CIM agent user identification**

After a CIM agent has been added to the **CIMOM** node, you can update the user ID, password, and other information.

To update the CIM agent user identification, follow this procedure:

1. In the Navigation Tree window, expand **Administrative Services > Data Sources**. Left-click **CIMOM**.
2. In the content pane, click the icon to the left of the agent for which you want to change information.
3. Information about the CIM agent is displayed in the Information window. The fields which can be updated are:

**Display Name**

The name of the CIM agent, as specified by the CIMOM provider, that will appear in the IBM Tivoli Storage Productivity Center interface.

**Description**

The optional description of the CIM agent that was entered on the Add CIMOM Agent window. Tivoli Storage Productivity Center

**Username**

The user ID used for authentication, if required by the CIM agent.

**Password**

The password for the CIM agent, if required by the CIM agent.

**Interoperability Namespace**

This namespace within the CIM agent allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how Tivoli Storage Productivity Center interacts with the CIM agent when retrieving information.

**Test CIMOM connectivity before updating**

Check box. Check this box to have Tivoli Storage Productivity Center ensure that communication to the CIM agent is working properly before updating information about that CIM agent.

4. Save your changes by clicking the **Save** icon.
  - If you have selected **Test CIMOM connectivity before updating**, the CIMOM connectivity check will be run when you save the changes.
  - If **Test CIMOM connectivity before updating** is not selected, a discovery will be started immediately after you save the changes.

**Testing a CIM agent connection**

Manually check to make sure that communication to the CIM agent is working properly.

To test a CIM agent connection, follow this procedure:

1. In the Navigation Tree, expand **Administrative Services > Data Sources**. Left-click **CIMOM**.

2. In the content pane, select a CIM agent. Click **Test CIMOM Connection**. a confirmation dialog appears: "Testing CIMOM connectivity can take up to several minutes in case of an incorrectly entered port number, network problems or an unpassed firewall. Would you like to continue anyway?" Select **yes** to perform the CIMOM connection test.
3. To close the Test CIMOM Connection window, click **OK**.

## Removing a CIM agent

You can remove a CIM agent from the navigation tree and the repository database. Data discovered by the CIM agent is not removed from the repository.

To remove a CIM agent, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services > Data Sources**. Click **CIMOM Agents**.
2. In the content pane, select a CIM agent. Click **Remove CIMOM**.
3. The Remove CIMOM confirmation window opens. Click **Yes**. The CIM agent is immediately deleted from the list.

## Show managed devices for a CIM agent

You can display the managed devices for a CIM agent.

To show the managed devices for a CIM agent, follow these steps:

1. In the Navigation Tree pane, expand **Administrative Services > Data Sources**. Click **CIMOM Agents**.
2. In the content pane, select a CIM agent. Click **Show Managed Devices**.
3. The CIMOM Managed Devices window opens displaying the managed devices. Click **OK** to close the window.

## Collecting CIM agent logs

You can collect logs for certain IBM CIM agents using the command line interface.

**Remember:** Storage systems that use the native interfaces (for example, DS8000, the XIV system, SAN Volume Controller, and Storwize V7000) do not use CIM agents.

1. Change to the directory where the CIM agent is installed.

|                                   |            |                                         |
|-----------------------------------|------------|-----------------------------------------|
| DS3000, DS4000, DS5000,<br>DS6000 | On Linux   | /opt/IBM/cimagent/cimom                 |
|                                   | On Windows | C:\Program Files\IBM\cimagent\<br>cimom |

2. Run one of the following commands:

|            |                 |
|------------|-----------------|
| On Linux   | collectLogs.sh  |
| On Windows | collectLogs.bat |

A **collectedLogs.zip** file is created.

**Attention:** this file is overwritten if you run the script again.

## Verifying that a CIM agent is running

You can verify that a CIM agent is running from the command line interface.

To verify that a CIM agent is up and running, run the following command:

```
telnet <IP> <port>
```

Where <IP> is the IP address of the system where the CIM agent is installed, and <port> is the port number. By default, this is 5989 for a secure connection and 5988 for an unsecure connection.

## Data agents and Storage Resource agents

This topic provides information about administering, configuring, and viewing information about Storage Resource agents and Data agents.

**Note:** In Tivoli Storage Productivity Center v4.2 and above, you can deploy Storage Resource agents only. If you want to install a Data agent, you must own a previous version of the product. See the documentation for the previous version of the product for information about how to install a Data agent.

See “Deployment considerations for Storage Resource agents” on page 120 for information about deploying Storage Resource agents.

### Viewing information about an agent

View detailed information about a Data agent or Storage Resource agent.

To view information about a Data agent or Storage Resource agent from the Tivoli Storage Productivity Center GUI, complete the following steps:

1. In the navigation tree, click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the content pane, click the magnifying glass icon to the left of a Data agent or Storage Resource agent under the list of agents. In the content pane, a notebook window opens. On the General page, the following information is displayed:

**Agent** Displays the name of an agent and indicates whether the agent is up or down.

**Port** Displays the port number on which an agent is listening for requests.

**Installation Location**

Displays where the agent is installed.

**Host Address**

Displays the host address of the computer on which an agent is running.

**Last update**

Displays the date and time when an agent was last updated.

**Time Zone**

Displays the current time zone for an agent.

**Connection Errors**

Displays any connection errors that occurred for an agent.

**Consecutive Errors**

Displays how many errors occurred consecutively.

**HBA Data**

Displays information about the HBA installed.

3. Click the **Details** tab to view more detailed information. The following information is displayed:

**Agent** Displays the name of the agent.

**Host** Displays the host name of the system on which the agent is running.



**Host ID**

Displays the identifier of the host computer in hexadecimal format. Tivoli Storage Productivity Center only collects Host ID information for Solaris systems. All other platforms will display GUID information in the Host ID column.

**Start Time**

Displays the date and time that the agent was started.

**Elapsed Time**

Displays the time elapsed since the agent was last updated.

**VM Size**

Displays the size (in MB) of the disk space that is allotted to virtual memory for a Data agent. This value does not apply to Storage Resource agents. It only applies to the size of the Java virtual machines for Data agents.

**Manufacturer**

Displays the manufacturer of the system.

**OS Type**

Displays the operating system that is running on the system.

**CPU Architecture**

Displays the architecture of the system.

**Job Count**

Displays the number of jobs scheduled to run against the agent.

4. Click the **Jobs** tab to view information about the jobs scheduled to run on the managed system.

**Viewing an agent log file**

The log file for a Data agent or Storage Resource agent contains informational, warning, and error messages for the previous sessions of the agent. You can use the log file to troubleshoot errors that might occur when starting an agent, processing, or shutting down the agent.

The log file is located in the agent installation directory:

**Storage Resource agents**

For Windows default: C:\Program Files\IBM\TPC\agent\log\  
<computer\_name>\agent.log.

For Linux or AIX (default): /opt/IBM/TPC/agent/log/  
<computer\_name>agent.log.

**Data agents**

For Windows (default): C:\Program Files\IBM\TPC\ca\subagents\TPC\  
Data\log\<computer\_name>.

For Linux or AIX (default): /opt/IBM/TPC/ca/subagents/TPC/Data/log/  
<computer\_name>.

For Data agents, by default, information from the last five sessions of the agent appears in the log window. A session starts when the agent is started and ends when the agent stops. You can increase or decrease how many sessions are kept in the history file by changing the value for the **logFilesKept** parameter in the agent.config file. The default location for the agent.config file is:

**Windows**

C:\Program Files\IBM\TPC\ca\subagents\TPC\Data\config\

## Linux or UNIX

/opt/IBM/TPC/ca/subagents/TPC/Data/config

To view the log for a Data agent or Storage Resource agent, complete the following steps:

1. In the navigation tree, click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the content pane, select a Data agent or Storage Resource agent and click **View Log**.

## Migrating a Data agent

Migrate a selected Data agent to a Storage Resource agent.

Tivoli Storage Productivity Center provides the following benefits when you migrate an existing Data agent or Inband Fabric agent to a Storage Resource agent:

- Storage Resource agents require less resources on a host computer than a Data agent or Fabric agent that is based on Common Agent Services. Additionally, you can deploy a Storage Resource agent as a daemon or non-daemon service.
- Reduced complexity when deploying agents. You can deploy Storage Resource agents directly from the Tivoli Storage Productivity Center user interface and they do not require you to install Agent Manager. For Data agents and Fabric agents, you must use the Tivoli Storage Productivity Center installation program and ensure that Agent Manager is registered with the Data server and Device server.
- Improved interface when deploying, upgrading, and administering agents. You can manage Storage Resource agents using the nodes in the **Administrative Services** section of the navigation tree. You can deploy and upgrade Storage Resource agents on schedules that you define.

To migrate a Storage Resource agent from the **Administrative Services > Data Sources** node in the user interface, complete the following steps:

1. In the navigation tree, expand **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the content pane, click **Migrate Agents**.
3. The Create Data/Fabric Agent Migration panel is displayed. Use this panel to schedule the migration of Data agents and Inband Fabric agents to Storage Resource agents.
4. Select **File > Save** to save the migration schedule. The agents on the computers that you specified are migrated according to the schedule that you defined.

The migration process deploys a Storage Resource agent on the target computer and then uninstalls the Data agent. The new Storage Resource agent retains the existing scheduled jobs that were defined for the Data agent. The scripts and batch reports for the Data agent are copied over to the Storage Resource agent. If for any reason the migration fails, Tivoli Storage Productivity Center restores the Data agent so that it can continue to monitor the target computer.

## Migrate Data/Fabric Agent

Define schedules to migrate existing Data agents or Inband Fabric agents to Storage Resource agents. The migration process deploys a Storage Resource agent on a target host and then uninstalls the existing agent.

Tivoli Storage Productivity Center provides the following benefits when you migrate an existing Data agent or Inband Fabric agent to a Storage Resource agent:

- Storage Resource agents require less resources on a host computer than a Data agent or Fabric agent that is based on Common Agent Services. Additionally, you can deploy a Storage Resource agent as a daemon or non-daemon service.
- Reduced complexity when deploying agents. You can deploy Storage Resource agents directly from the Tivoli Storage Productivity Center user interface and they do not require you to install Agent Manager. For Data agents and Fabric agents, you must use the Tivoli Storage Productivity Center installation program and ensure that Agent Manager is registered with the Data server and Device server.
- Improved interface when deploying, upgrading, and administering agents. You can manage Storage Resource agents using the nodes in the **Administrative Services** section of the navigation tree. You can deploy and upgrade Storage Resource agents on schedules that you define.

To migrate a Data agent or Fabric agent to a Storage Resource agent from the **Administrative Services > Configuration** node in the user interface, complete the following steps:

1. Use either of the following methods to access the panel for scheduling an agent migration:
  - Expand **Administrative Services > Data Sources** in the navigation tree. Left-click **Data/Storage Resource Agents**. In the content pane, click **Migrate Agents**.
  - Expand **Administrative Services > Configuration** in the navigation tree. Right-click **Migrate Data/Fabric Agent** and select **Create Data/Fabric Agent Migration** from the menu.

The content pane displays a panel for defining a migration schedule.

2. Use the following pages on the panel to define information about the migration:

#### **Computers**

Use this page to select the computers and groups of computers that contain the Data and Fabric agents that you want to migrate to Storage Resource agents.

#### **When to Run**

Use this page to determine when the Data and Fabric agents in a schedule are migrated.

#### **Options**

Use this page to specify the runtime settings of a Storage Resource agent, the method for migrating an agent, and the login credentials for an agent. If you include multiple agents in a migration schedule, the settings that you define on this panel are applied to all those agent migrations.

**Alerts** Use this page to define the actions that occur if a Data/Fabric Agent Migration schedule fails.

3. Select **File > Save** to save the migration schedule. The agents on the computers that you specified are migrated according to the schedule that you defined.

The migration process deploys a Storage Resource agent on the target computer and then uninstalls the Data agent. The new Storage Resource agent retains the existing scheduled jobs that were defined for the Data agent. The scripts and batch reports for the Data agent are copied over to the Storage Resource agent. If for any reason the migration fails, Tivoli Storage Productivity Center restores the Data agent so that it can continue to monitor the target computer.

## Upgrading the Storage Resource agent

This topic provides information for upgrading a downlevel Storage Resource agent to the current release.

To upgrade a Storage Resource agent, follow these steps:

1. In the navigation tree, expand **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the content pane, select one or more agents for which you want to upgrade. Click **Upgrade Agents**.
3. The Create Storage Resource Agent Upgrade panel is displayed. Use this panel to schedule an upgrade of the Storage Resource agent.

**Note:** The Storage Resource agent upgrade job delivers the software upgrade packages to the agent computer. The job log displays the status of the delivery. The actual status of the upgrade is found in the agent log on the agent computer. If the agent log indicates that the upgrade failed and the state of the Storage Resource agent remains in the "Upgrading agent software" status, try restarting the agent and running the upgrade job again.

## Storage Resource Agent Upgrades

You can configure IBM Tivoli Storage Productivity Center for a Storage Resource agent upgrade.

When you apply maintenance to Tivoli Storage Productivity Center, you can upgrade the Storage Resource agent now or upgrade the agent at a later time. To ensure that all your agents are at the current release level and to manage your network load, schedule upgrades on a regular basis.

To configure Tivoli Storage Productivity Center for a Storage Resource agent upgrade, complete the following steps:

1. In the navigation tree, expand **Administrative Services > Configuration > Storage Resource Agent Upgrades**.
2. Right-click **Storage Resource Agent Upgrades**. Click **Create Storage Resource Agent Upgrade**.
3. The Create Storage Resource Agent Upgrade panel is displayed in the content pane. You have the following tabs:

### Computers

You can select Computer Groups or Computers.

### When to Run

You can specify how often to run the job or how to handle time zones.

### Options

You can specify to force an upgrade for an agent, maintenance upgrade for an agent, or language pack to apply to an agent.

**Alerts** You can specify the triggering conditions and triggering actions.

Enter your information in the appropriate fields.

4. Save and run the job.

## Configuring a Data agent

You can change the configuration of a Data agent by modifying the values that are specified in the **agent.config** file. The configuration parameters that are set in the

**agent.config** file include the host name of the system on which the Data server is installed, the port on which the Data agent listens, and log file options.

To modify the configuration parameters for a Data agent, complete the following steps:

1. Using a text editor, modify the **agent.config** file. This file is located in one of the following directories:

|                   |                                              |
|-------------------|----------------------------------------------|
| <b>On AIX</b>     | /opt/IBM/TPC/ca/subagents/TPC/Data/config    |
| <b>On Linux</b>   | /opt/IBM/TPC/ca/subagents/TPC/Data/config    |
| <b>On Windows</b> | <directory>\TPC\ca\subagents\TPC\Data\config |

where <directory> is the directory where the Data agent is installed.

2. Save the modified **agent.config** file.
3. From the IBM Tivoli Storage Productivity Center GUI, in the navigation tree, click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
4. In the content pane, select a Data agent.
5. Click **Read Config** from the **Select Action** menu. The Data server reads the **agent.config** file that you modified in step 1 and applies the changes.

### Configuring tracing for an agent

You can configure and enable tracing for a Data agent or Storage Resource agent. This provides extensive logging of exception and trace messages; it can provide useful information for the IBM Support center.

You must have administrator authority to perform this procedure.

To configure tracing for a Data agent or a Storage Resource agent, complete the following steps:

1. In the navigation tree, click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the content pane, select a Data agent or Storage Resource agent. Under the **Select Action** menu, click **Configure Tracing**. The Tracing Configuration window opens.
3. Select the **Enable Trace** check box.
4. Specify the tracing options:
  - a. In the **Level** field, select one of the following options:
    - DEBUG\_MIN**  
Minimum level of debugging
    - DEBUG\_MID**  
Medium level of debugging
    - DEBUG\_MAX**  
Maximum level of debugging
  - b. In the **Maximum Number of Files** field, specify the maximum number of trace files that are created. When this number is reached, the oldest file is overwritten. By default, this is set to five.
  - c. In the **Maximum File Size (kilobytes)** field, specify the maximum size of each trace file.
5. Click **OK**. The settings are saved to the **AgentTraceLog.config** file

## Configuring a Data agent

You can change the configuration of a Data agent by modifying the values that are specified in the **agent.config** file. The configuration parameters that are set in the **agent.config** file include the host name of the system on which the Data server is installed, the port on which the Data agent listens, and log file options.

To modify the configuration parameters for a Data agent, complete the following steps:

1. Using a text editor, modify the **agent.config** file. This file is located in one of the following directories:

|                   |                                              |
|-------------------|----------------------------------------------|
| <b>On AIX</b>     | /opt/IBM/TPC/ca/subagents/TPC/Data/config    |
| <b>On Linux</b>   | /opt/IBM/TPC/ca/subagents/TPC/Data/config    |
| <b>On Windows</b> | <directory>\TPC\ca\subagents\TPC\Data\config |

where <directory> is the directory where the Data agent is installed.

2. Save the modified **agent.config** file.
3. From the IBM Tivoli Storage Productivity Center GUI, in the navigation tree, click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
4. In the content pane, select a Data agent.
5. Click **Read Config** from the **Select Action** menu. The Data server reads the **agent.config** file that you modified in step 1 and applies the changes.

## Reading the configuration file (Data agent)

If you make changes to the agent.config file and want to verify those changes, you can read the configuration file.

The agent.config file contains parameters that you can change to customize the behavior of an agent. The default directory for the agent.config file is:

### Windows

C:\Program Files\IBM\TPC\ca\subagents\TPC\Data\config\

### Linux or UNIX

/opt/IBM/TPC/ca/subagents/TPC/Data/config

To verify the contents of the configuration file, complete the following steps:

1. In the navigation tree, click **Administrative Services > Data Sources > Data/Storage Resource agents**.
2. In the right pane, select a Data agent. Click **Read Config**.
3. The configuration file is displayed.

## Disabling an agent

You can disable a Data agent or a Storage Resource agent. Disabling the agent signifies that the Data agent or Storage Resource agent is unavailable and the Data server should not contact that agent for any job processing.

You might want to disable a Data agent or Storage Resource agent in the following situations:

- The agent computer is undergoing maintenance and will be unavailable. This prevents the server from flagging the agent as "down" if it cannot reach the agent. The amount of times that the server tries to contact the agent is defined by the **agentErrorLimit** parameter in the server.config file.

- The agent computer is busy with resource-intensive processing and you do not want to add any jobs to that processing load.

While disabled, the agent name in the user interface is outlined with a red circle with a line through it. The server does not attempt to contact the agent.

To disable a Data agent or Storage Resource agent, complete the following steps:

1. In the navigation tree, click **Administrative Services > Data Sources > .** Left-click **Data/Storage Resource Agents**.
2. In the content pane, select a Data agent or Storage Resource agent. Under the Select Action menu, click **Disable**.

## Enabling an agent

You can enable an agent that is currently disabled. After it is enabled, the Data server resumes communication with that agent. This action is available for disabled agents only.

If the server cannot contact an agent, the server automatically flags the agent as "down". Click **Enable** to reestablish communication between the agent and server. The number of times the server tries to contact the agent is specified in the **agentErrorLimit** parameter in the **server.config** file. The default is 3.

The default directory for the **server.config** file is:

### Windows

C:\Program Files\IBM\TPC\Data\config

### Linux or UNIX

/opt/IBM/TPC/Data/config

**Tip:** Clicking **Check** also enables an agent that is disabled.

To enable an agent, complete the following steps:

1. In the navigation tree, click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the content pane, select an agent. Under the Select Action menu, click **Enable**.

## Shutting down an agent

You can shut down Data agents and Storage Resource agents (run as daemon processes) from the graphical user interface.

To shut down an agent, complete the following steps:

1. In the navigation tree, click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the content pane, select a Data agent or Storage Resource agent. Under the **Select Action** menu, click **Shutdown**. This button is unavailable for Storage Resource agents that are deployed as non-daemons.
3. Select how you want an agent to shut down. You can select from the following options:

### Normal

Shut down the agent software and let all running processes complete. The agent continues to accept new jobs that are submitted and is not shut down until the following conditions are met:

- All running processes are complete.

- There are no new jobs submitted.

**Immediate**

Shut down the agent software when the currently running processes complete. No new jobs are accepted by the agent and shutdown occurs immediately after the last job completes.

**Abort** Shut down the agent software and stop whatever processes are currently running.

You cannot start a Data agent from the Data/Storage Resource agent panel in the IBM Tivoli Storage Productivity Center user interface. To start a Data agent, you must access the computer on which the agent is installed and start it manually.

For Storage Resource agents that are run as daemon processes, you can use the **Start** button on the Data/Storage Resource agent panel to start those agents. See “Starting a Storage Resource agent” for more information about how to start Storage Resource agents.

## Starting a Storage Resource agent

You can start a Storage Resource agent if the daemon process for that Storage Resource agent is down.

You can only perform this task for Storage Resource agents running as a daemon service.

To start a Storage Resource agent, complete the following steps:

1. In the navigation tree, click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the content pane, select a Storage Resource agent. Under the **Select Action** menu, click **Start**.
3. Enter the user ID and password that Tivoli Storage Productivity Center uses when logging in to the host computer for a Storage Resource agent and click **Ok**.

## Checking whether an agent can communicate with the Data Server

You can check whether an agent can communicate with the Data server.

To check an agent, complete the following steps:

1. In the navigation tree, click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the content pane, select an agent and click **Check**.

If a Data agent is selected, this operation verifies that the Data server can communicate with the selected agent.

If a Storage Resource agent is selected, this operation verifies that the Data server can communicate with the host system where the Storage Resource agent is deployed with the appropriate authentication information. After verifying that the Data server can communicate with the selected agent, the process checks for the existence of the runtime files for Storage Resource agents.

## Changing authentication for a Storage Resource agent

There are times when you might need to change the user ID, password, and certificate location for a Storage Resource agent.



To change the authentication for a Storage Resource agent, follow these steps:

1. In the navigation tree, expand **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the content pane, select one or more Storage Resource agents for which you want to change authentication information. Use ctrl + click or shift + click to select multiple agents.
3. Click **Change Authentication** from the **Select Action** menu.
4. In the Change Authentication window, enter the following information:

**User** Enter the user ID that IBM Tivoli Storage Productivity Center should use when logging on to the computer on which the Storage Resource agent is installed.

**Requirement:** This ID must have administrative or root privileges when installing a Storage Resource agent on a computer.

**Password**

Enter the password.

**Certificate Location**

Enter the fully qualified path of the certificate on the computer where the Data server is located. This file is used for certificate-based authentication. If you do not enter a value in this field the default location is used. When deploying Storage Resource agents, Tivoli Storage Productivity Center uses the Secure Shell (SSH) protocol to communicate with the target computer.

**Passphrase**

Enter the passphrase for the certificate file. The passphrase was created when the certificate was generated.

Click **Save**.

**Related Tasks:**

See **Planning > General Planning > Planning for Storage Resource agents** section in the *IBM Tivoli Storage Productivity Center Installation and Configuration Guide* for additional configuration information about certificates used after deployment.

## **Adding Storage Resource agents**

You can add Storage Resource agents to your system by scheduling a job or by adding an agent immediately.

To add a Storage Resource agent, complete the following steps:

1. In the navigation tree, click **Administrative Services > Configuration**. Right-click **Storage Resource Agent Deployments**. On the menu, click **Create Storage Resource Agent Deployments**.
2. The Create Storage Resource agent Deployments window opens displaying the Computers page. Enter the following information:

**Header fields**

**Creator**

Predefined creator name.

**Name** Predefined name.

**Description (optional)**

Enter a user-defined description name for the job.

### Enabled

Check this box to enable Storage Resource agent deployments.

### Add Host List

Click this button to manually enter names and logon credentials for the computers on which you want to deploy Storage Resource agents.

If you click this button, the Login Information window opens. You can enter information in the following ways:

- Click **Add Agents from MS Directory** to install Storage Resource agents on one or more Windows computers that are members of a Windows domain. Enter the domain controller name and login information about the displayed window. On the next window, select the domain and click **Get List of Domain Computers**. The host names of computers that you add using this method are displayed in the **Remote Agent Machines** section of the window. In that section, you must also enter an installation location for each computer and indicate whether you are reinstalling a Storage Resource agent on a computer.
- Click **Get Agent List From file** to retrieve a list of computer names from a flat text file. The computers that are listed in the file are automatically added to the deployment job and displayed in the **Remote Agent Machines** section of this window. In that section, you must also enter an installation location for each computer and indicate whether you are reinstalling a Storage Resource agent on a computer.

The file that you retrieve must contain one host name or IP address per line. For example:

```
systemxyz.storage.usa.ibm.com
9.79.179.179
systemzyx.storage.usa.ibm.com
9.89.189.189
```

If Tivoli Storage Productivity Center discovers any syntax problems in the file, none of the host entries in the file are added to the deployment job.

- Enter the computer names or IP addresses in the Remote Agent Machines table. Enter the installation locations. If you do not enter a value in this field the default location is used. The default installation directories are:
  - For Windows: C:\Program Files\IBM\TPC\agent
  - For Linux and UNIX: /opt/IBM/TPC/agent
- Select **Force** under the following circumstances:
  - If an earlier Storage Resource agent installation failed and there are damaged agent files on the computer that cause further installations to fail. If you select this option, Tivoli Storage Productivity Center attempts to overwrite the previous failed deployment on the computer with a new Storage Resource agent.
  - If you want an existing Storage Resource agent to communicate with an additional Tivoli Storage Productivity Center server. To do this, you must create the deployment job from the additional Tivoli Storage Productivity Center server to which you want the Storage Resource agent to communicate.

To install the Storage Resource agent to communicate with two different servers using the same location, follow the example in these steps:

- a. From server 1 (S1), install the Storage Resource agent to the location C:\Program Files\IBM\TPC\agent using port 9510.
- b. From server 2 (S2), install the Storage Resource agent in the same location using the same port 9510 and specify the **Force** option. This option is required when you install the Storage Resource agent in the same location.

If you want to install two Storage Resource agents on the same system, each with a different port, install the agents using different port numbers and different locations.

**Note:**

- You cannot change the communication type for a Storage Resource agent (daemon or non-daemon) when you select the **Force** option. Make sure to select the same communication type as the existing Storage Resource agent when you create a deployment job.
- If a Storage Resource agent exists on a target computer and you do not select **Force**, an error occurs during validation and the Storage Resource agent is not installed.
- Select the **Initial Scan** option to perform a file system scan on the computer where a Storage Resource agent is deployed. If this box is checked, the scan is run immediately after a Storage Resource agent is deployed.
- Click the **Windows Services Information** button to enter login information that Tivoli Storage Productivity Center uses to create a Windows service for the Storage Resource agent. Use the Windows Service Information window to select the local system account, enter information about an existing user account, or create a new local account. If you do not specify a Windows service account, the default itcauser account is created. This action is only available for agents that are deployed to Windows computers as daemon services. This button is disabled if Use Daemon Service for Runtime Operation is not selected.

Enter the following information:

**User** Enter the user ID that Tivoli Storage Productivity Center uses when logging in to the host computer to install a Storage Resource agent. The value in the field is applied to all the computers that appear in the **Remote Agent Machines** table.

**Requirement:** This ID must have administrative or root privileges on the target computer.

**Password**

Enter the password for the user ID.

**Certificate Location**

Enter the fully qualified path of the certificate on the computer where the Data server is located. This file is used for certificate-based authentication. If you do not enter a value in this field the default location is used.

**Passphrase**

Enter the passphrase for the certificate file. The passphrase was created when the certificate was generated.

**Port** Enter the port number on which a Storage Resource agent listens for requests. The default is 9510.

**Tip:** This value is required if you run the Storage Resource agent as a daemon service.

**Use Daemon Service for Runtime Operation**

Select this option to run a Storage Resource agent on the monitored computer as a daemon service.

See Storage Resource agent Deployment Considerations for more information about deploying agents as daemon and non-daemon services.

**Restriction:** Do not select this option to run the Storage Resource agent as a stand-alone executable file on the monitored computer.

**Validate before Save**

Check this box to indicate that Tivoli Storage Productivity Center should attempt to communicate with the computers that are added to the deployment job when you click the **Save** button on this page. If the validation fails for a computer, the agent is not installed on that computer.

**Edit Selected Entries**

Click this button to edit settings for the selected Storage Resource agents. This information includes user ID, password, certificate location, passphrase, and port. You must submit the deployment job for the changes to take effect.

**Remove**

Click this button to remove a Storage Resource agent from the deployment job.

**Tip:** This button does not remove the agent from Tivoli Storage Productivity Center if it is already installed; it only removes it from the deployment job.

3. Click the **When to run** tab to enter the following information:

**How often to run**

Specify a time to run:

- **Run now**
- **Run once at** (specify a date and time to run)

**How to handle time zones**

Specify a time zone to use:

- **Use the time zone that the server runs in**
- **Use this time zone** (select a time zone)

4. Click the **Alert** tab to specify the following information:

**Triggering-Condition**

The triggering condition you can specify is: Storage Resource agent Deployment Failed

### Triggered-Actions

Choose from the following check boxes:

- **SNMP Trap**
- **TEC / OMNibus Event**
- **Login Notification**
- **Windows Event Log**
- **Email**

Depending on what action you select, you might have other choices to make. For example, if you select the **Windows Event Log** check box, the **Event Type** field becomes active so that you can specify a severity for the event in the Windows event log.

5. Click **File > Save**.

### Related Tasks:

See the *Planning > General Planning > Planning for Storage Resource agents* section in the *Tivoli Storage Productivity Center Installation and Configuration Guide* for additional configuration information about certificates used after deployment.

Check the `TPCD_#####.log` file in the Tivoli Storage Productivity Center installation directory to view the error messages related to failed validations. The following are the default locations for this log file:

- Windows: `C:\Program Files\IBM\TPC\data\log`
- UNIX/Linux: `/opt/IBM/TPC/data/log`

If multiple Storage Resource agents are set up to probe or scan the same storage resources, the Storage Resource agent that was added to Tivoli Storage Productivity Center first is used for the probe or scan. Therefore, only data that is gathered by the first Storage Resource agent is shown.

### Collecting service data

Collect service data about the selected Storage Resource agent (v4.2 or later only). Service data includes diagnostic information such as logs, trace files, configuration information, and computer details. Use this data to troubleshoot any errors that might occur during startup, processing, or shutdown of a Storage Resource agent.

If the collection of service data is successful, Tivoli Storage Productivity Center displays a message that indicates where the resulting archive file (.zip) is stored. The archive file is stored in a directory on the same computer as the Tivoli Storage Productivity Center server. The file is located in the following default locations:

- Windows: `C:\Program Files\IBM\TPC\data\log\SRATraces\  
SRA_computer_name\TPCServiceInfo.zip`
- UNIX or Linux: `/opt/IBM/TPC/data/log/SRATraces/SRA_computer_name/  
TPCServiceInfo.zip`

`SRA_computer_name` represents the name of the computer on which the Storage Resource agent is deployed.

If the collection of service data fails, Tivoli Storage Productivity Center displays a message that indicates the job was unsuccessful. See the server log file or the services script output for more information about why the data collection failed. These files are located in the following default locations:

- Server log file (on the computer where the Tivoli Storage Productivity Center server is installed):
  - Windows: c:\Program Files\IBM\TPC\data\log
  - UNIX or Linux: /opt/IBM/TPC/data/log
- Services script output file (on the computer where the Storage Resource agent is deployed):
  - Windows: C:\Program Files\IBM\TPC\agent\service\TPCServiceInfo.html
  - UNIX or Linux: /opt/IBM/TPC/agent/service/TPCServiceInfo.html

where *SRA\_computer\_name* represents the name of the computer on which the Storage Resource agent is located.

To collect service data for a Storage Resource agent, complete the following steps:

1. In the navigation tree, click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource agents**.
2. In the content pane, select a Storage Resource agent.
3. In the **Select Action** menu, click **Collect Service Data**.

### Enabling or disabling scripts for Storage Resource agents

You can enable or disable the running of scripts that are sent from the Tivoli Storage Productivity Center server to a Storage Resource agent. If enabled, a Storage Resource agent accepts and runs scripts that are sent from the server for scheduling. If disabled, a Storage Resource agent does not accept scripts that are sent from the server and only runs scripts that are stored on the computer where the agent is deployed.

To enable or disable the running of scripts, complete the following steps:

1. In the navigation tree, expand **Administrative Services > Data Sources > Data/Storage Resource Agents**.
2. In the content pane under the Select Action menu, click **Enable Scripts** to enable the running of scripts. Click **Disable Scripts** to disable the running of scripts.

### Enabling or disabling the monitoring of fabrics by the Storage Resource agents

Enable or disable the monitoring of fabrics by the Storage Resource agent. By default a Storage Resource agent is enabled for fabric monitoring.

To enable or disable the monitoring of fabrics by the Storage Resource agent, complete the following steps:

1. In the navigation tree, expand **Administrative Services > Data Sources > Data/Storage Resource Agents**.
2. In the content pane under the Select Action menu, click **Enable Fabric** to enable the monitoring of fabrics. Click **Disable Fabric** to disable the monitoring of fabrics.

### Help command for Storage Resource agent

The **help** command for the Storage Resource agent provides information about the parameters for installing, uninstalling, and upgrading Storage Resource agents.

For information about the Storage Resource agent commands, run the **help** command. Follow these steps:

1. Go to the installation location for the Storage Resource agent:

```

cd <CD_installation_location>
2. Run the following command:
bin/Agent -help
3. The output from the help command is as follows:
Usage:
Agent -INSTALL
 [-COMMTYPE DAEMON -AGENTPORT portnumber]
 [-FORCE]
 -INSTALLLOC pathname
 -SERVERIP address[,address,...]
 -SERVERPORT portnumber
 [-USERID username -PASSWORD password -CERT file -PASSPHRASE phrase]

Agent -UNINSTALL
 [-FORCE]
 -SERVERNAME servername

Agent -UPGRADE
 -INSTALLLOC pathname

```

For information about how to install the Storage Resource agent using commands, see the information about installing agents in the *Tivoli Storage Productivity Center Installation and Configuration Guide*.

For information about how to uninstall the Storage Resource agent using commands, see the information about uninstalling agents in the *Tivoli Storage Productivity Center Installation and Configuration Guide*.

For information about how to upgrade the Storage Resource agent using commands, see the information about upgrading agents in the *Tivoli Storage Productivity Center Installation and Configuration Guide*.

## Enabling and disabling NFS timeouts

You can enable or disable NFS timeouts for a Storage Resource agent when running scans on network file system (NFS) mounts. Enabling an agent for NFS timeouts prevents scans from locking up indefinitely when NFS mounts the agents are monitoring become unavailable.

This option determines the number of seconds a Storage Resource agent waits for a status system call on an NFS file system before it times out. The default wait time is 10 seconds.

**Note:** Enabling this option might decrease the performance of scans slightly. Disabling this option might increase the performance of scans slightly.

You might want to enable NFS timeouts when the NFS server is down or an NFS file system is stale. This condition can also occur when a scan is running for an agent and that agent is uninstalled.

The **FileSystemScanNFSTimeout** option is in the Data server configuration file (`server.config`). The default location of the `server.config` file is `c:\Program Files\IBM\TPC\config` for Windows and `/opt/IBM/TPC/config` for UNIX and Linux.

To enable or disable NFS timeouts, follow these steps:

1. In the navigation tree, click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.

2. In the content pane, select a Storage Resource agent.
3. To enable NFS timeouts for a Storage Resource agent, under the Select Action menu, click **Enable NFS Timeout**.
4. To disable NFS timeouts for a Storage Resource agent, under the Select Action menu, click **Disable NFS Timeout**.

### **Deleting or uninstalling an agent using the GUI**

You can delete or uninstall a Storage Resource agent and all the data that is collected by that agent from the database repository. The information that is collected by the agent is no longer available within Tivoli Storage Productivity Center reports. When you delete an agent, the agent is uninstalled and you can no longer activate the agent.

To delete or uninstall a Storage Resource agent, complete the following steps:

1. In the navigation tree, click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the content pane, select an agent and under the Select Action menu, click **Delete**.
3. In the confirmation dialog, select **Yes Delete the agent**.

### **Enabling or disabling automatic upgrade for agents**

This topic provides information for enabling or disabling the automatic upgrade function for a Storage Resource agent.

When you install or upgrade Tivoli Storage Productivity Center, the **Enable Auto Upgrade** function is turned off.

**Note:** You cannot upgrade a Data agent to a Storage Resource agent. If you want to migrate a Data agent to a Storage Resource agent, use the **Administrative Services > Configuration > Migrate Data/Fabric Agents** node in the user interface.

To enable the program so that a Storage Resource agent can be automatically upgraded, follow these steps:

1. In the navigation tree, expand **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the content pane, select one or more agents for which you want an automatic upgrade performed.
3. Under the Select Action list box, click **Enable Auto Upgrade** (this action displays as **Disable Auto Upgrade** if the Auto Upgrade function is on).

### **Exporting a Storage Resource agent list**

You can export a Storage Resource agent list to a formatted flat file. This list is used to manage user authentication through a script file. This option is enabled for Storage Resource agents that use non-daemon based communication only.

This list contains the host name, installation location, user ID, encrypted password, and fully qualified certificate file for the selected Storage Resource agents. You can change the user authentication data such as user ID, password, and fully qualified certificate location.

When the file is created, the passwords for the agents are written as "\*\*\*\*\*" for security reasons. Before using this file with the **Update Storage Resource Agent**



**List** action, you must enter an encrypted or non-encrypted password for each host. See the description of the **Update Storage Resource Agent List** action for more information about passwords.

Storage Resource agents can either be run as a daemon service or non-daemon service (on-demand service). If the agent is run as a non-daemon service, there is no agent process running on the managed host. Programs on the host system are run as needed.

To export a Storage Resource agent list, complete the following steps:

1. In the navigation tree, click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the content pane, select a Storage Resource agent and click **Export Storage Resource Agent List** from the **Select Action** menu.

### Updating a Storage Resource agent list

The Storage Resource agent list file is used to manage user authentication through a script file. This list contains the host name, installation location, user ID, encrypted password, and fully qualified certificate file for the selected Storage Resource agents. The password is in encrypted format with @ENC@ appended to the encrypted password. You can change the user authentication data such as user ID, password, and fully-qualified certificate location.

To update a Storage Resource agent list, complete the following steps:

1. In the navigation tree, click **Administrative Services > Data Sources**. Left-click **Data/Storage Resource Agents**.
2. In the content pane, select the Storage Resource agents you want to update.
3. Under the Select Action menu, click **Update Storage Resource Agent List** to update information about the agent.
4. Edit the information in that file to match the changes that you want to make for the agent. You can update the logon credentials for the Storage Resource agents that are listed with the information in a file that you specify. The file must use the following format for the host names you want to update:

```
host_name|installation_location|user_ID|password
```

If the Data server cannot find any Storage Resource agent with a host name and location that you specify in the file, that entry is ignored and the error is logged.

Keep in mind the following considerations when including a password in the update file:

- For non-encrypted passwords: If you use "|" in the line for an agent, you must provide an "escape" character before it. The escape character is "\". For example, if the password you want to include in an update file is "native|agent", use native\|agent in the update file.

If you use "\" in the line for an agent, you must provide an additional "\" as an escape character. For example, if the password is "native\agent", use native\\agent in the update file.

- For encrypted passwords:
  - To encrypt a password, start IBM Tivoli Storage Productivity Center's command-line tool (tpctool) from the \cli directory where Tivoli Storage Productivity Center is installed and run the following command: tpctool encrypt [password], where [password] represents the password you want to encrypt. This command will generate a text string that represents the

encrypted version of the password. For example, output might appear like the following: iDroqToC070ubh5i4mxMHQ==. Copy this encrypted password to the correct location in the update file.

- To include an encrypted password in an update file, you must add the following after the encrypted text: @ENC@. For example, an encrypted password in an update file might appear like this:  
iDroqToC070ubh5i4mxMHQ==@ENC@

**Note:**

- Do not include @ENC@ at the end of non-encrypted passwords.
- See the *IBM Tivoli Storage Productivity Center Command-Line Interface Reference* for more information about **tpctool**.

## Registering the Storage Resource agent with a different server

This topic provides information on how to register the Storage Resource agent with a different server.

To register the Storage Resource agent with a different server, complete the following steps:

1. You have a Storage Resource agent that is installed on Server A and you want the Storage Resource agent to also point to Server B.
2. From Server B, in the navigation tree, expand **Administrative Services > Configuration > Storage Resource Agent Deployments**. Right-click **Storage Resource Agent Deployments**.
3. Specify the same port and location with the **force** option.

**Note:** For an RXA Storage Resource agent, just specify the same location with the **force** option.

4. Run the deployment job. The Storage Resource agent now points to Server A and Server B. If you no longer want the Storage Resource agent to point to Server A, just go to the GUI on Server A and delete the Storage Resource agent.

## Managing the daemon Storage Resource agent service

Use this information to start and stop the daemon Storage Resource agent service for the Virtual I/O Server.

### Starting and stopping the daemon Storage Resource agent service

To start or stop the daemon service, follow these steps:

1. Log in to the Virtual I/O Server by using the **padmin** user ID.
2. Run the following command to set up the AIX environment:  
oem\_setup\_env
3. Change to the directory where the Storage Resource agent is located. For example:
  - To stop the service, run this command:  
/SRA\_install\_directory/agent/bin/agent.sh stop
  - To start the service, run this command:  
/SRA\_install\_directory/agent/bin/agent.sh start

## Registering the Data agent with a different server

This topic provides information on how to register the Data agent with a different server.

Use the PROBE\_ME file to get an existing Data agent to register with a different IBM Tivoli Storage Productivity Center Data server from the one that it currently uses. This file is also useful to reestablish the connection between the Data server and Data agent (for example, in the case of a Data server reinstallation). To change the registration of the Data agent with a different server, complete the following steps:

1. Change the name of the server in the agent.config file by performing the following steps:
  - a. In the <TPC\_install>/ca/subagents/TPC/Data/config directory, open the agent.config file.
  - b. Change the serverHost property to the IP address of the new Tivoli Storage Productivity Center server.
2. Create a blank text file called PROBE\_ME (all in uppercase letters with no txt extension) in the root of the agent's Data agent directory (<TPC\_install\_dir>/ca/subagents/TPC/Data/).
3. Restart the agent.

This will initiate a probe job to run and the agent to register with the server specified in the agent.config file. If you put any text in the PROBE\_ME file, a default scan will also be run.

### **Deployment considerations for Storage Resource agents**

There are a number of guidelines and limitations that you must consider when managing the IBM Tivoli Storage Productivity Center agents in your environment.

Use the following information when deploying Storage Resource agents and migrating existing Data agents:

- Multiple Storage Resource agents that are probing or scanning the same storage resources
- Platforms that support Storage Resource agents deployments
- Product functions that are not available for storage devices monitored by Storage Resource agents
- Required authority for deploying Storage Resource agents
- Orphan zones
- Firewalls and Storage Resource agent deployments
- Multiple computers in a deployment
- Communication between the Tivoli Storage Productivity Center server and a Storage Resource agent
- Daemon and non-daemon services
- Authentication between the Tivoli Storage Productivity Center server and a Storage Resource agent
- Storage Resource agents and Data agents on the same computer
- Time zones for computers monitored by Storage Resource agents
- UNIX systems using RSH protocol
- Deployments on Windows - NetBIOS setting
- Deployments on Windows - User Account Control (UAC) remote restrictions
- Additional information about Storage Resource agents

#### **Multiple Storage Resource agents that are probing or scanning the same storage resources**

If multiple Storage Resource agents are set up to probe or scan the same storage resources, the Storage Resource agent that was added to Tivoli

Storage Productivity Center first is used for the probe or scan. Therefore, only data that is gathered by the first Storage Resource agent is shown.

#### **Platforms that support Storage Resource agents for the same**

For a list of platforms on which you can deploy Storage Resource agents, go to the Tivoli Storage Productivity Center support site at [http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli\\_Storage\\_Productivity\\_Center\\_Standard\\_Edition](http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center_Standard_Edition), click the **Documentation** link, and enter **Platform Support: Agents, Servers and GUI** in the **Search** box. Click the link to the document for the appropriate release.

#### **Product functions that are not available for storage devices monitored by Storage Resource agents**

Before you migrate an existing Data agent to a Storage Resource agent or deploy a new Storage Resource agent, make sure that the product functions you want to use on the monitored devices are available for those agents. The following functions are not available for storage devices that are monitored by Storage Resource agents:

- NetWare monitoring. You must use Data agents to monitor NetWare Servers.
- Monitoring of certain relational databases. For list of relational databases that can be monitored by Storage Resource agents, go to the Tivoli Storage Productivity Center support site at [http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli\\_Storage\\_Productivity\\_Center\\_Standard\\_Edition](http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center_Standard_Edition), click the **Documentation** link, and enter **Platform Support: Agents, Servers and GUI** in the **Search** box. Click the link to the document for the appropriate release.
- Changing zone configuration and reporting of HBA, fabric topology, or zoning information for fabrics connected to hosts running Linux on IBM Power Systems or zSeries hardware. These limitations also apply to Storage Resource agents on all guest operating systems for VMware configurations.

If you are already using these functions in Tivoli Storage Productivity Center, it is recommended that you do not migrate the corresponding Data agents to Storage Resource agent. If you plan to use these functions in the future, you must install Data agents on the computers that you want to monitor.

#### **Required authorities for deploying and running Storage Resource agents**

Before you can create deployment schedules and deploy Storage Resource agents on target computers, you must meet the following requirements:

- To create deployment schedules, you must be logged in to Tivoli Storage Productivity Center with a user ID that has the **Productivity Center Administrator** role. See the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* for more information about user roles.
- To deploy Storage Resource agents on target computers, you must provide a user ID that has administrative rights on those computers. You enter this ID when creating a deployment schedule. Tivoli Storage Productivity Center uses this ID to log in to the target computers and install and configure the necessary runtime files for the agents.

The user under which a Storage Resource agent (daemon or non-daemon) runs must have the following authorities on the target computers:

- On Linux or UNIX, the user must have root authority. By default, an agent runs under the user 'root'.
- On Windows, the user must have Administrator authority and be a member of the Administrators group. By default, a Storage Resource agent runs under the 'Local System' account.

### Orphan zones

Storage Resource agents do not collect information about orphan zones. An orphan zone is a zone that does not belong to at least one zoneset.

### Firewalls and Storage Resource agent deployments

Before you can deploy a Storage Resource agent on a computer, you must turn off the firewall on that computer. If you do not turn off the firewall, the deployment fails. To turn off the firewall on a Windows 2008 computer, complete the following steps:

1. On the desktop of the Windows 2008 computer, click **Start > Administrative Tools**.
2. Click **Windows Firewall with Advanced Security**.
3. Click **Windows Firewall Properties**.
4. Change the **Firewall state** field to **Off** on the following tabs: **Domain Profile, Private Profile, Public Profile**.
5. Click **OK** to accept the changes and exit.
6. Deploy a Storage Resource agent to the Windows 2008 computer.

### Multiple computers in a deployment

You can define a deployment schedule to include multiple computers on which to install Storage Resource agents. The computers you include in a deployment schedule must share the same administrative user ID and password. IBM Tivoli Storage Productivity Center uses these user credentials to log in to the computers when installing Storage Resource agents. If the computers in a deployment schedule do not share the same administrative user credentials, you must create separate deployment jobs for them.

**Note:** When you deploy a Storage Resource agent to multiple computers, a globally unique identifier (GUID) is created for each computer (if one does not exist).

### Communication between the Tivoli Storage Productivity Center server and a Storage Resource agent

The Tivoli Storage Productivity Center server connects to a monitored computer when a Storage Resource agent is deployed and whenever a data collection schedule runs against that agent.

During deployment, the server communicates with the target computer using one of the following protocols:

- Windows server message block protocol (SMB)
- Secure Shell protocol (SSH)
- Remote execution protocol (REXEC)
- Remote shell protocol (RSH)

After deployment, the type of communication between the server and agent on that computer depends on whether you deployed the agent as daemon service or non-daemon service. See the *IBM Tivoli Storage*

*Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* for more information about these protocols.

### **Daemon and non-daemon services**

You can deploy a Storage Resource agent as a daemon or non-daemon service:

- A Storage Resource agent that is deployed as a daemon service runs in the background on the monitored computer and listens for requests from the Tivoli Storage Productivity Center server. Connectivity between the server and agent is established using SSL. The server and agent have their respective certificates and no additional information is required besides those certificates and the security that is provided by the SSL protocol.
- A Storage Resource agent deployed as a service on demand (non-daemon service) runs as a stand-alone executable file on the monitored computer. Communication from the server to the agent uses the same protocol that was used during the deployment of the agent. Communication from the agent to the server uses SSL.

### **Authentication between the Tivoli Storage Productivity Center server and a Storage Resource agent**

Tivoli Storage Productivity Center requires the correct authentication information (user ID, password, port, certificate location, or passphrase) for monitored computers each time it communicates with Storage Resource agents on those computers. If the authentication information changes for a host computer on which a Storage Resource agent is deployed, the authentication information for that agent must be updated using the **Change Authentication** or **Update Storage Resource agent List** functions on the Data/Storage Resource agents panel located under the **Administrative Services > Data Sources** node in the navigation tree.

### **Storage Resource agents and Data agents on the same computer**

You cannot deploy a Storage Resource agent on a computer where a Data agent is already installed *and* pointing to the same Data server as the Storage Resource agent. You can deploy a Storage Resource agent on the same computer as a Data agent if those agents communicate with different Data servers and use different ports when listening for requests.

### **Time zones for computers monitored by Storage Resource agents**

The time zones of computers that are monitored by Storage Resource agents are shown as Greenwich Mean Time (GMT) offsets in Tivoli Storage Productivity Center reports. The time zones of computers that are monitored by Data agents are shown as the local time zone in Tivoli Storage Productivity Center reports. For example, a computer located in Los Angeles shows the following time zones in the By Computer report in Asset reporting.

- If monitored by a Storage Resource agent: (GMT-8:00) GMT-8:00
- If monitored by a Data agent: (GMT-8:00) America/Los\_Angeles Time zone

### **UNIX systems using Remote Shell protocol (RSH)**

If RSH is configured to use a user ID and password, the connection fails. To successfully connect to a system using RSH, you must set up the `.rhosts` file (in the home directory of the account). RSH must be configured to accept a login from the system that is running your application.

## Deployments on Windows - NetBIOS setting

If you want to install a Storage Resource agent or Data agent on Windows targets, the **Enable NetBIOS over TCP/IP** option must be selected in the Control Panel settings for the computer's network connections properties. To set this option, click **Start > Settings > Control Panel > Network and Dial-Up Connections > <some\_connection> > Properties > Internet Protocol (TCP/IP) > Advanced > WINS > Enable NetBIOS over TCP/IP**.

See the documentation for your firewall to determine if these ports are not blocked for inbound requests.

To determine if security policies are blocking the connection ports, click **Start > Settings > Control Panel > Administrative Tools**.

Depending on whether your policies are stored locally or in Active Directory, follow these directions:

### Locally stored policies

For locally stored policies, click **Start > Settings > Control Panel > Administrative Tools > Local Security Policy > IP Security Policies on Local Computer**.

### Policies stored in Active Directory

For policies stored in Active Directory, examine the IP security policies and edit or remove filters that block the ports:

- Click **Administrative Tools > Default Domain Security Settings > IP Security Policies on Active Directory**.
- Click **Administrative Tools > Default Domain Controller Security Settings > IP Security Policies on Active Directory**.

For all Windows systems, the Server service must be running to connect to a Windows system using the Windows protocol.

The following table lists the ports reserved for NetBIOS. Ensure that these ports are not blocked.

| Port | Description                                          |
|------|------------------------------------------------------|
| 135  | NetBIOS Remote procedure call. (Not currently used.) |
| 137  | NetBIOS name service.                                |
| 138  | NetBIOS datagram. (Not currently used.)              |
| 139  | NetBIOS session (for file and print sharing).        |
| 445  | CIFS (on Windows XP).                                |

For Windows Server 2008, shares must be shared for the Guest or Everyone accounts, and password protected sharing must be disabled. To disable password protected sharing, follow these steps:

1. Click **Control Panel > Networking and Sharing Center**.
2. Click the down arrow next to **Password protected sharing**.
3. Click **Turn off password protected sharing**.
4. Click **Apply**. Exit from the Control Panel.

## Deployments on Windows 2008 - User Account Control (UAC) remote restrictions

If you are planning to install Storage Resource agents remotely on a Windows 2008 system, you must disable the User Account Control (UAC) remote restrictions on the Windows system. User Account Control is a security component on Windows.

**Note:** This task contains steps that tell you how to modify the registry. Serious problems might occur if you modify the registry incorrectly. Therefore, make sure that you follow these steps carefully. For added protection, back up the registry before you modify it. Then you can restore the registry if problems occur. For information about how to back up and restore the registry, see <http://support.microsoft.com/kb/322756/>.

To disable UAC remote restrictions, follow these steps:

1. Click **Start > Run**. Enter **regedit** and click **OK**.
2. Locate and click the following registry subkey:  
HKEY\_LOCAL\_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\  
Policies\System
3. If the LocalAccountTokenFilterPolicy registry entry does not exist, follow these steps:
  - a. On the **Edit** menu, click **New > DWORD Value**.
  - b. Enter **LocalAccountTokenFilterPolicy** as the name for the DWORD value in the content pane. Click **Enter**.
  - c. Right-click **LocalAccountTokenFilterPolicy**, then click **Modify**.
  - d. In the Edit DWORD Value window, enter **1**, then click **OK**.  
This value can be 0 or 1:
    - 0 This value builds a filtered token. This value is the default value. The administrator credentials are removed.
    - 1 This value builds an elevated token.
  - e. Exit the registry editor.

#### **Additional information about Storage Resource agents**

See *Planning for Storage Resource agents* in the *IBM Tivoli Storage Productivity Center Installation and Configuration Guide* for more information about Storage Resource agents.

## **In-band Fabric agents**

In-band Fabric agents are agents that are used by Fabric Manager for in-band SAN operations.

**Note:** In Tivoli Storage Productivity Center v4.2 and above, you can deploy Storage Resource agents only. If you want to install an Inband fabric agent (Fabric agent), you must own a previous version of the product. See the documentation for the previous version of the product for information about how to install a Fabric agent.

For information about the supported agent types for switch performance management and fabric zone configuration, see “Supported Fabric Agent Types” on page 1143.

### **Displaying information about an in-band Fabric agent**

You can view information about an in-band Fabric agent including the status, IP address, host name, operating system, and Host Bus Adapter data.

To display information about an in-band Fabric agent, follow this procedure:

1. In the Navigation tree pane, expand **Administrative Services > Data Sources**. Left-click **Inband Fabric**.



2. In the content pane, click the icon to the left of the Fabric agent for which you want information for. In the content pane, a notebook window opens. On the General page, the following information is displayed:

**Status** The status of the agent.

**IP Address**

The IP address of the agent.

**Host name**

The host name of the agent.

**Operating System and Version**

The operating system and version of the agent (for example, Windows 9.1: Service Pack 2).

**HBA Data**

The Host Bus Adapter data for the agent.

## Checking the in-band agent connection

You can check if an in-band Fabric agent is running and whether the version of the agent is up to date. A check also enables an agent that is disabled.

To check an in-band Fabric agent connection, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources**. Left-click **Inband Fabric**.
2. In the content pane, select an agent and click **Check**.

**Note:** The Fabric agent status is taken from the database. Because the agent does not send indications when it is shut down, the status may not be current. The `setup.properties` file for the Fabric agent has the following default parameter set: `monitorInterval=10`. This indicates that the monitor will be refreshed every 10 minutes. You can change the monitor setting for a different refresh time.

The status of the agent is displayed in the Agent Status window. If updates are required, they are automatically installed.

## Viewing an agent log

Use the content of the log file to troubleshoot any errors that might occur during startup, processing or shutdown. Service alerts, if any, are displayed in red.

To display the log of an in-band Fabric agent, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources**. Left-click **Inband Fabric**.
2. In the content pane, select an agent and select **View Log**. In the content pane, the agent log is displayed.

## Deleting an in-band Fabric agent

You can delete an in-band Fabric agent from the navigation tree and the database. Data discovered by the agent is removed from the database repository. This action is available for inactive agents only. You must first stop the agent before deleting the information for the agent. The agent is not uninstalled from IBM Tivoli Storage Productivity Center.

To remove information for an in-band Fabric agent, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources**. Left-click **Inband Fabric Agents**.

2. In the content pane, select the agent and click **Delete**.

## Migrating an in-band Fabric agent

Migrate a selected in-band Fabric agent to a Storage Resource agent.

The migration process deploys a Storage Resource agent on the target computer and then uninstalls the in-band Fabric agent. Before the migration, Tivoli Storage Productivity Center must shut down the Fabric agent. The agent is not shut down until 1) all running processes are complete, and 2) there are no new schedules submitted for the agent.

## Out-of-band Fabric agents

Out-of-band agents are used by Fabric Manager for out-of-band SAN operations.

For information about the supported agent types for switch performance management and fabric zone configuration, see “Supported Fabric Agent Types” on page 1143.

### Displaying information about an out-of-band agent

You can view information about an out-of-band agent including the IP address, status, and SNMP community.

To display information about an out-of-band agent, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources**. Click **Out of Band Fabric Agents**.
2. In the content pane, click the icon to the left of the agent for which you want information. In the content pane, a notebook window opens. On the General Information page, the following information is displayed:

**Status** The status of the agent.

#### Host Name

The host name of the system on which the agent is installed.

#### IP Address

The IP address of the system on which the agent is installed. This column displays Internet Protocol Version 4 (IPv4) and Internet Protocol Version 6 (IPv6) addresses as appropriate.

#### SNMP Community Read

The name of the Simple Network Management Protocol (SNMP) community to which the agent belongs. The SNMP community name acts as a password that is shared by one or more SNMP hosts. The community name is used to authenticate messages being received by this SNMP host. This field is optional and might be blank if the SNMP community has not been set.

**Note:** The default SNMP community is public. If the community name is not the correct community name for your environment, the out-of-band agent might not be able to properly perform scans.

#### SNMP Community Write

The community string for SNMP writes. Tivoli Storage Productivity Center uses the write community string to refresh information for the Out of Band Fabric Agent. The default is private. This value is used for Cisco switches only.

## Updating out-of-band agent information

After an out-of-band agent has been added, you can update the SNMP community information.

To update information for an out-of-band agent, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services > Data Sources**. Left-click **Out of Band Fabric Agents**.
2. In the content pane, click the icon to the left of the agent.
3. In the content pane, a notebook window opens. On the General page, update the following information:

### SNMP Community Read

Enter the name of the Simple Network Management Protocol (SNMP) community to which the agent belongs. The SNMP community name acts as a password that is shared by one or more SNMP hosts. The community name is used to authenticate messages being received by this SNMP host. This field is optional.

**Note:** The default SNMP community is public. If this is not the correct community name for your environment, the out-of-band agent might not be able to properly perform scans.

### SNMP Community Write

Enter the community string for SNMP writes. Tivoli Storage Productivity Center uses the write community string to refresh information for the Out of Band Fabric Agent. The default is private. This value is used for Cisco switches only.

## Checking the out-of-band agent connection

You can check to determine if an out-of-band agent is running.

To check the agent connection, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services > Data Sources**. Left-click **Out of Band Fabric Agents**.
2. In the content pane, select an agent and click **Check**. The status of the agent is displayed in the Agent Status window. If updates are required, they are automatically installed.

## Adding an out-of-band Fabric agent

This topic provides information about adding an out-of-band Fabric agent.

You can add an out-of-band Fabric agent with the Configure Devices wizard or by using the manual steps in this procedure. To add an out-of-band Fabric agent with the Configure Devices wizard, see “Configure Devices” on page 305.

To add an out-of-band Fabric agent manually, complete the following steps:

1. In the navigation tree, expand **Administrative Services > Data Sources**. Left-click **Out of Band Fabric Agents**.
2. In the content pane, click **Add**.
3. In the Add Out of Band Fabric Agent window, complete the host name, SNMP community read (optional), and SNMP community write (optional) fields. Click **OK**.

## Deleting an out-of-band agent

You can remove an out-of-band agent from the navigation tree and the database. Data discovered by the agent is not removed from the database repository.

To remove an out-of-band agent, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources**. Left-click **Out of Band Fabric**.
2. In the content pane, select an agent and click **Delete**.

## Removing out-of-band Fabric agent data

If you have a McDATA Intrepid 10000 director and have out-of-band Fabric agents as well as in-band or SMI-S agents, and you have run an out-of-band fabric discovery job, you must delete the data collected from the out-of-band Fabric agent because that data is invalid. This topic provides information on how to delete the out-of-band agent data.

To remove out-of-band Fabric agent data, complete the following steps:

1. Open the IBM Tivoli Storage Productivity Center GUI.
2. Expand **Administrative Services > Data Sources**. Left-click **Out of Band Fabric**.
3. In the content pane, select the out-of-band agent and click **Delete**.

## Tivoli Storage Productivity Center servers

Use the **Administrative Services > Data Sources > IBM Tivoli Storage Productivity Center Servers** function to manage the relationships between a *master* IBM Tivoli Storage Productivity Center server and its *subordinate* Tivoli Storage Productivity Center servers.

Creating master and subordinate relationships between Tivoli Storage Productivity Center servers enables you to use a single interface to generate reports based on data and metrics collected by multiple servers in a storage environment:

- A **master** server is a server that performs normal monitoring and reporting of storage resources like a standard server, but also gathers the storage information (using Tivoli Storage Productivity Center server probes) that has been collected by subordinate servers.
- A **subordinate** server is a server that monitors and reports on storage resources like a standard server, but also communicates with the master server during Tivoli Storage Productivity Center server probes. During these probes, the master server collects the storage information gathered by a subordinate server's agents and stores that information in its own database repository.

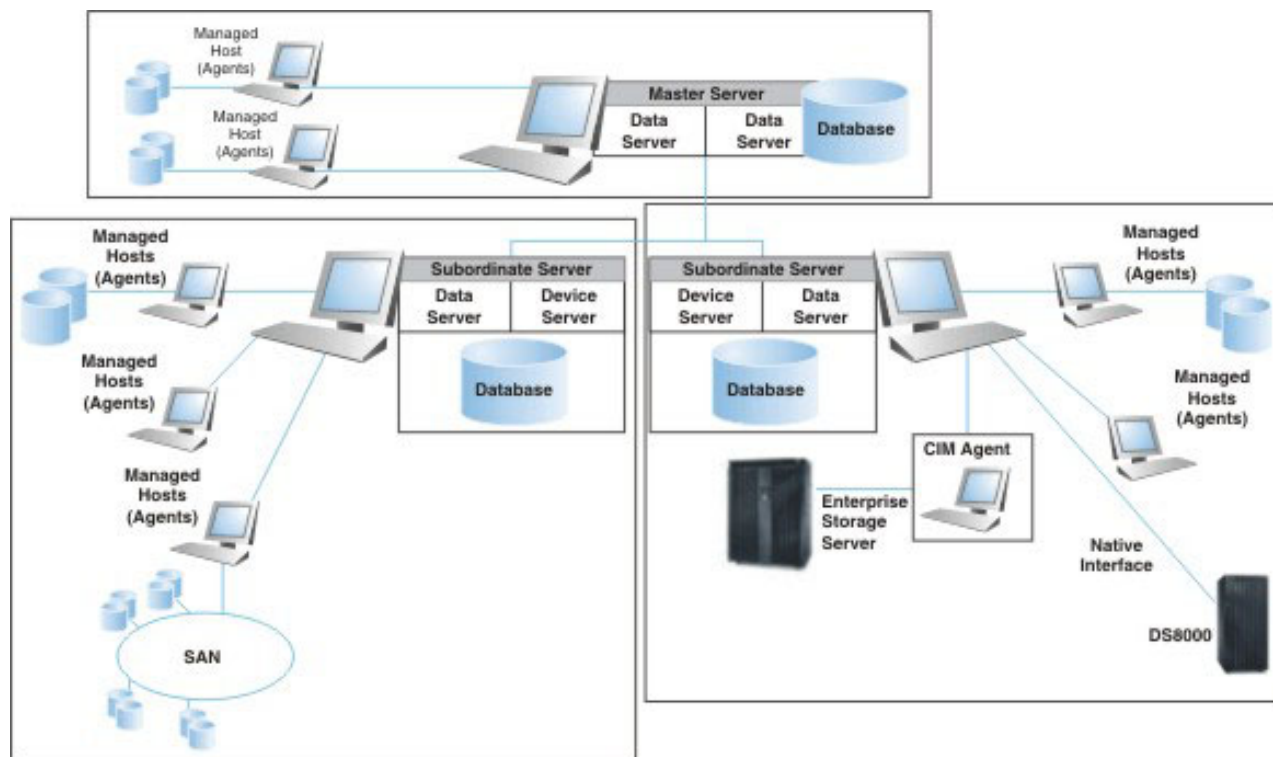


Figure 21. Master and subordinate server architecture

The rollup reports that reflect the storage information collected by the master server from subordinate servers are available in the **Tivoli Storage Productivity Center > Rollup Reports** node of the master server's navigation tree.

If the master server is located on an IPv6-only server, it can communicate with existing subordinate servers under the following conditions:

- The subordinate servers are upgraded to Tivoli Storage Productivity Center V4.1 or higher and
- The IPv6 protocol is enabled on the machines where they are located.

Before you can configure and manage subordinate servers, keep in mind the following:

- The master server must be up and running.
- You must be logged in to the user interface as an Tivoli Storage Productivity Center administrator or superuser
- We recommend that the master server should monitor no more than 500 unique data sources. This number includes subordinate servers, Storage Resource agents, in-band Fabric agents, CIMOM agents, and VM servers (VMWare).
- We recommend that each subordinate server monitors no more than 1200 unique data sources. This number includes Storage Resource agents, in-band Fabric agents, CIMOM agents, and VM servers (VMWare). Once this threshold has been met for a server, a new server should be deployed and all new agents pointed to it.
- After upgrading to Tivoli Storage Productivity Center v4.2 or later from a previous version, you must run probes against a master server's monitored storage assets to have information about those assets appear in **Tivoli Storage Productivity Center > Reporting > Rollup Reports**.

- When you remove a data source (agent) from a subordinate server or a storage entity that it monitors, the master server continues to display information about that agent or storage resource in rollup reports for a specified amount of time. You can specify how quickly that information is removed by accessing the Administrative Services > Configuration > Removed Resource Retention panel and changing the settings for **TPC servers**. Additionally, if a master server is unable to contact a subordinate server for the number of days defined in the **TPC Servers** setting on the Removed Resource Retention panel, that server is removed from the subordinate server as well as information about the storage entities it was monitoring.
- A host might appear more than once in a rollup report and its data is counted multiple times in report totals. Multiple entries for a host might occur if the following conditions are true:
  - An agent is removed from a host, the GUID (globally unique identifier) for the host is changed, and then an agent is redeployed on the host.
  - The amount of time that information about a host is retained is longer than the time it took to remove the original agent and change the GUID for the host.

For example, if the number of days to retain information about a host is 14 days, and then you 1) remove the agent on the host, 2) change the GUID of the host, and 3) redeploy an agent on that host all in less than 14 days, duplicate entries for the host appears in rollup reports. When the 14 days have passed, the entry for the host before the GUID change is removed from the database and only one entry appears for the host in rollup reports.

You can specify how quickly information is removed by changing the settings for **TPC servers** on the Administrative Services > Configuration > Removed Resource Retention panel.

## Adding an Tivoli Storage Productivity Center Server

Learn how to add an Tivoli Storage Productivity Center server as a subordinate sever.

Keep in mind the following when adding a Tivoli Storage Productivity Center server as a subordinate server:

- How subordinate servers are configured to monitor storage entities within your environment can determine if fragmentation occurs within rollup reports. See Rollup reports - Fragmentation for more information.
- Tivoli Storage Productivity Center servers must be at version 3.3 or greater to be added as a subordinate server.
- We recommend that the master server should monitor no more than 500 unique data sources. This number includes subordinate servers, Storage Resource agents, in-band Fabric agents, CIMOM agents, and VM servers (VMWare).
- We recommend that each subordinate server monitors no more than 1200 unique data sources. This number includes Storage Resource agents, in-band Fabric agents, CIMOM agents, and VM servers (VMWare). Once this threshold has been met for a server, a new server should be deployed and all new agents pointed to it.
- If the master server is located on an IPv6-only server, it can communicate with existing subordinate servers under the following conditions:
  - The subordinate servers are upgraded to Tivoli Storage Productivity Center V4.1 or higher and
  - The IPv6 protocol is enabled on the machines where they are located.

To add a server as a subordinate server, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services > Data Sources > Tivoli Storage Productivity Center Servers**.
2. In the content pane, click **Add Tivoli Storage Productivity Center Server**.
3. In the Add Tivoli Storage Productivity Center Server window, enter the following information:

**Host Name**

Enter the host name of the machine on which the subordinate server is installed.

**Host Device server Port**

Enter the port with which the subordinate server's Device server is listening for requests. The default port for Tivoli Storage Productivity Center servers is 9550.

**Host Authentication Password**

Enter the password for the subordinate Tivoli Storage Productivity Center server. This host authentication password is entered when installing Tivoli Storage Productivity Center. Contact your system administrator or the person who installed Tivoli Storage Productivity Center to determine this password.

**Display Name (optional)**

Enter a name for the subordinate server.

**Description (optional)**

Enter a description for the subordinate server.

**Test Tivoli Storage Productivity Center Server Connectivity before adding**

Select this check box to have the master server connect and log into the subordinate server using the information provided on this window. This test is performed when you click **Save**.

4. Click **Save** to add the server as a subordinate server.
  - If you selected **Test Tivoli Storage Productivity Center Server Connectivity before adding**, IBM Tivoli Storage Productivity Center attempts to connect to the subordinate server using the information provided on this window.
  - If you did not select **Test Tivoli Storage Productivity Center Server Connectivity before adding**, the subordinate server is automatically added as a subordinate server but has a connection status of **Unknown** on the **Administrative Services > Data Sources > Tivoli Storage Productivity Center Servers** window.

**Note:** If the login information for a subordinate server is incorrect, an error occurs and an associated error message dialog box is displayed.

## Viewing subordinate server details

Learn how to view information about the subordinate servers that have been associated with a master server.

After you add a subordinate server to a master server, basic information about that server appears on the **Administrative Services > Data Sources > TPC Servers** window. To view additional information about a subordinate server on the TPC Servers window, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources > TPC Servers**.

2. In the content pane, double-click a server row or click the icon next to a server for more detailed information about that server. Information about a subordinate server appears on the **Administrative Services > Data Sources > TPC Servers > TPC Server Details** window.

## Testing the connection to subordinate servers

Learn how to test the connection from a master server to its associated subordinate servers.

After you add a subordinate server to a master server, basic information about that server appears on the **Administrative Services > Data Sources > TPC Servers** window. To verify that the master server can use the information provided to connect to subordinate server, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources > TPC Servers**.
2. In the content pane, highlight the row for a subordinate server and click **Test TPC Server Connection** to test its connection to the master server. A message window will appear with the results of the tested connection. The information that appears in the **Connection Status** field will be updated to reflect its status.

## Removing subordinate servers

Learn how to remove a subordinate server from its association with a master server.

Use the **Administrative Services > Data Sources > TPC Servers** window to remove a subordinate server from a master server. Removing a subordinate server removes subordinate server information from:

- The master server database repository
- All rollup reports
- The **TPC Servers** list on the **Administrative Services > Data Sources > TPC Servers** window

To remove a subordinate server, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources > TPC Servers**.
2. In the content pane, highlight the row of a subordinate server and click **Remove TPC Server**. A confirmation window appears.
3. Click **Yes** to confirm the removal of the subordinate server.

## Modifying subordinate server information

Learn how to modify information about the subordinate servers that have been associated with a master server.

After you add a subordinate server to a master server, basic information about that server appears on the **Administrative Services > Data Sources > TPC Servers** window. You can edit the information for a subordinate server by modifying its Display Name, Description, and Host Authentication Password. To view modify information about a subordinate server on the TPC Servers window, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources > TPC Servers**.



2. In the content pane, double-click a server row or click the icon next to a server to edit information about that server. Information about a subordinate server appears on **Administrative Services > Data Sources > TPC Servers > TPC Server Details**.

3. Edit the information in the following fields:

**Host Authentication Password**

Enter the password for the subordinate IBM Tivoli Storage Productivity Center server.

**Display Name (optional)**

Enter a name for the subordinate server.

**Description (optional)**

Enter a description for the subordinate server.

**Test TPC Server Connectivity before adding**

Select this check box to have the master server connect and log into the subordinate server using the information provided on this window. This test is performed when you click **Save**.

4. Select **File > Save** . If successful, the TPC Servers window appears and any new values for **Display Name** and **Description** are shown for a subordinate server.

## Rollup reports - fragmentation

Fragmentation occurs when related storage entities are being monitored by different subordinate IBM Tivoli Storage Productivity Center servers.

The following example shows when fragmentation can occur using subordinate servers and rollup reports:

- **Example:** Host A is assigned a volume from storage subsystem 1. Host A is being monitored only by one Tivoli Storage Productivity Center subordinate server (Server 1) and storage subsystem 1 is being monitored by another subordinate server (Server 2). Server 1 and Server 2 are subordinate servers to the master server (Server 3).

**Reason for fragmentation:** In this example, the host and subsystem are fragmented because they are not being monitored by the same Tivoli Storage Productivity Center server. This will cause the **Correlated Volume Space** value in the rollup reports for storage subsystems and volumes to be incorrect.

To prevent fragmentation when using master and subordinate server relationships to generate rollup reports, we recommend the following:

- Do not monitor a storage subsystem and the hosts to which its volumes are assigned with different subordinate servers that are associated with the same master server.
- Do not monitor a fabric and its related switches with different subordinate servers that are associated with the same master server.
- If an agent machine is using storage from a subsystem, then both the agent and subsystem should be monitored by the same subordinate server.
- Configure all the subordinate servers to discover all the storage subsystems within your environment. Note that if a subordinate server is monitoring a subsystem with an agent other than the agent that is using storage on that subsystem, the value for the **Correlated Volume Space** column in rollup reports will be incorrect. For example:
  - Host A has storage from Subsystem 1
  - Subordinate server B is Monitoring Host A and Subsystem 1

- Subordinate server C is Monitoring Subsystem 1 and probed the subsystem more recently than server B.
- The master server is monitoring subordinate servers B and server C.

The master server will report the storage subsystem data from server C. Because the master server does not know about Host A, the value for the **Correlated Volume Space** column will be incorrect.

## Storage systems

This topic provides information about administering, configuring, and viewing information about monitored storage systems in your IBM Tivoli Storage Productivity Center environment.

You can add many different types of storage systems to your Tivoli Storage Productivity Center environment. Use the native storage system interfaces provided to connect to the following storage systems:

- DS8000
- SAN Volume Controller
- Storwize V7000
- IBM Storwize V7000 Unified
- The XIV system

For other storage systems, use the Common Interface Model (CIM) agent interfaces.

### Adding a storage system

Learn how to add a storage system to your IBM Tivoli Storage Productivity Center environment.

Tivoli Storage Productivity Center provides multiple ways to add a storage system to your environment. Use the Configure Devices wizard to guide you through the steps for adding a storage subsystem using its native interface or a CIMOM, or manually add a CIMOM and discover the storage subsystem that it manages. See “Configure storage subsystems” on page 306 for details of how to use the Configure Devices wizard.

To add a storage subsystem, use any of the following methods:

- If you want to add a new storage subsystem to Tivoli Storage Productivity Center using its native interface, use the Configure Devices wizard. The wizard guides you through the steps for adding a storage subsystem and running a discovery. To use this method, you can click the Configure Devices icon in the toolbar or you can complete the following steps:
  1. In the navigation tree, expand **Administrative Services > Data Sources > Storage Subsystems**.
  2. Click **Add**. This action launches the Configure Devices wizard. See “Configure storage subsystems” on page 306 for details of how to use the Configure Devices wizard.
- If you want to add a CIMOM to Tivoli Storage Productivity Center and discover the storage subsystem that it manages, use the Configure Devices wizard. The wizard guides you through the steps for adding a CIMOM and running a discovery. To use this method, you can click the Configure Devices icon in the toolbar or you can complete the following steps:
  1. In the navigation tree, expand **Administrative Services > Data Sources > Storage Subsystems**.

2. Click **Add**. This action launches the Configure Devices wizard. See “Configure storage subsystems” on page 306 for details of how to use the Configure Devices wizard.
- If you want to manually add a new CIMOM to Tivoli Storage Productivity Center and then manually discover the storage subsystem that it manages, use the following steps:
  1. In the navigation tree, expand **Administrative Services > Data Sources** and click **CIMOM Agents**.
  2. In the content pane, click **Add CIMOM**. The Add CIMOM window opens. See “Adding a CIM agent connection” on page 244 for details of how to add the CIMOM.
  3. To discover the storage subsystem that a CIMOM manages, in the navigation tree pane, expand **Administrative Services > Discovery** and click **CIMOM**. See “Discover storage subsystem, tape library, and fabric information with CIM agents” on page 330 for details about how to discover the storage subsystems managed by the CIMOMs that you manually added.

## Viewing storage system details

Learn how to view information about monitored storage systems in your IBM Tivoli Storage Productivity Center.

After you add storage systems, basic information about those systems appears on the **Administrative Services > Data Sources > Storage Subsystems** panel. To view information about a storage system on the Storage Subsystems panel, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources > Storage Subsystems**.
2. In the content pane, you can view basic information about the storage systems in the Storage Subsystems area.

### Magnifying glass icon

The function of this icon depends on the value that is in the Credential Status column:

- If the value is **Present**, click the icon to view and modify more detailed information about the system.
- If the value is **Missing**, click the icon to start the Configure Devices wizard and provide the configuration information for the storage system.
- If the value is **Required**, click the icon to open the subsystem credentials update dialog box.

**Name** The name of the storage system.

**Label** A user-defined name of the system.

**Type** The type of storage system.

### IP Address


The IP address of the storage system.

### Credential Status

The status of the credentials that are used to connect to the storage system. Possible values are:

- **Present** (the credentials are available).
- **Missing** (no credentials are available because the system had been discovered using the automated discovery process. To add the

required credentials, click the  icon to open the Configure Devices wizard. Follow the steps in the wizard.

- **Required** (the system must be updated from CIM object manager [CIMOM] to native interface connection). To update the connection method, click the  icon to open the Update Subsystem Credentials dialog box.

#### Connection Status

The status of the connection between Tivoli Storage Productivity Center and the storage system.

#### Data Source IP Address

The IP address of the entity that Tivoli Storage Productivity Center uses to connect to the storage system. If Tivoli Storage Productivity Center uses the native interface of a system for connection, this IP address is the same as the address in the **IP Address** field. If Tivoli Storage Productivity Center uses a CIMOM for connection, the data source IP address is the address for the CIMOM.

## Modifying XIV system details

Learn how to modify information for an IBM XIV Storage System device that connects to IBM Tivoli Storage Productivity Center.

After you add the XIV system device, basic information about that system appears on the **Administrative Services > Data Sources > Storage Subsystems** panel. To modify information about an XIV system device on the Storage Subsystems panel, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources > Storage Subsystems**.
2. In the content panel, double-click a storage subsystem row or click the icon next to a row to edit information about that storage subsystem. Information about the storage subsystem appears on the Storage Subsystems Details window.
3. You can edit the information in the following fields:

#### Username

Displays the user name for the device.

#### Password

Displays the password for the device.

The information in the following fields cannot be edited:

#### Storage Subsystem

Displays the name of the device.

#### IP Address

Displays the IP address or host name for the device.

4. Select **File > Save**. If successful, the Storage Subsystems window appears and any new values for the storage subsystem are shown.

## Modifying SAN Volume Controller or Storwize V7000 storage system details

Learn how to modify information about SAN Volume Controller or Storwize V7000 storage systems that connect to IBM Tivoli Storage Productivity Center.

After you add a SAN Volume Controller or Storwize V7000 device, basic information about that system appears on the **Administrative Services > Data**

**Sources > Storage Subsystems** panel. To modify information about a SAN Volume Controller on the Storage Subsystems panel, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources > Storage Subsystems**.
2. In the content panel, double-click a storage subsystem row or click the icon next to a row to edit information about that storage subsystem. Information about the storage subsystem appears on the Storage Subsystems Details window.
3. You can edit the information in the fields that are specified in bold in the window:

**Software Version**

Displays the version of the SAN Volume Controller or Storwize V7000 device.

**Storage Subsystem**

Displays the name of the device.

**IP Address**

Displays the IP address or host name for the device.

**Select Key**

Lists the actions that are available for the public Secure Shell (SSH) key:

- **Use existing uploaded key.** This action is the default. If you select this action, the **Admin Username** and **Admin Password** fields are optional.
- **Upload new key.** If you select this action, the **Admin Username** and **Admin Password** are required.

**Admin Username**

If you selected **Upload new key** in the **Select Key** list, enter the name of a user that has the administrator role for the SAN Volume Controller or Storwize V7000 cluster that contains the device. If you selected **Use existing uploaded key**, this field is optional.

**Admin Password**

Enter the password for the user that you entered in the **Admin Username** field.

**Username**

Enter an authorized user for the device. This field is active if you have selected **Upload new key** in the **Select Key** list and have entered a user name and password in the **Admin Username** and **Admin Password** fields. To search for a user, click **Select User**. You must enter a valid administrator name and password to retrieve a list of users.

**Private SSH Key**

Displays the path for the private SSH key. If you want to change this path, enter a new path or click the **Browse** button to search for a new path.

**Passphrase (optional)**

Displays the passphrase for the SSH key pair.

4. Select **File > Save**. If successful, the Storage Subsystems window appears and any new values for the storage subsystem are shown.

## Modifying Storwize V7000 Unified storage system details

After you add an Storwize V7000 Unified storage system, information about the system is displayed the Storage Subsystems Details page. You can modify the information that is in the available fields that are on this page.

Storwize V7000 Unified is composed of Storwize V7000 and Storwize V7000 File Module. You can connect to only Storwize V7000 or to both Storwize V7000 and Storwize V7000 File Module. Storwize V7000 provides block-level data and Storwize V7000 File Module provides file-level data.

To modify information about a Storwize V7000 Unified storage system, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources > Storage Subsystems**.
2. On the Storage Subsystems page, double-click a storage subsystem row to edit information about that storage subsystem. The following information about the storage subsystem is displayed:

### Storwize V7000

#### Storage Subsystem

Displays the name of the storage system.

#### IP Address

Enter the IP address or host name of the Storwize V7000 storage system. You can enter an IPv4 or IPv6 address, for example:

- 127.0.0.1
- [2001:DB8:0:0:0:0:0:0/32]

#### Select Key

An SSH key is required to access Storwize V7000. Select one of the following actions for the SSH key:

- **Upload new key:** Select this option to upload a new key. Public key content is extracted from the private key that is provided in the **Private SSH Key** field and is uploaded to the Storwize V7000 cluster. If you select this action, the **Administrator User Name** and **Administrator Password** fields are required.
- **Use existing uploaded key:** Select this option if you want to use a public key that was already uploaded by using a method other than through Tivoli Storage Productivity Center, such as the Storwize V7000 web interface. Enter the location of the private key file that corresponds to the public key in the **Private SSH Key** field. The **Administrator User Name** and **Administrator Password** fields are not required if you select this option.

#### Administrator User Name

If you selected **Upload new key** in the **Select Key** list, enter the name of a user that has the Security Administrator role for the Storwize V7000 cluster that contains the storage system. Tivoli Storage Productivity Center uses this value during the setup process to contact the cluster and configure the SSH keys for the user. This user name must belong to the Security

Administrator role, which has privileges to modify other user accounts. Otherwise, Tivoli Storage Productivity Center cannot configure the SSH keys.

If you selected **Use existing uploaded key**, this field is optional.

#### **Administrator Password**

Enter the password for the user name that you entered in the **Administrator User Name** field.

#### **User Name**

Enter an authorized user name for the storage system. Tivoli Storage Productivity Center uses this value to monitor and configure the storage system. If the user name does not exist, it is created. This user name must belong to the Administrator role.

This field is active if you selected **Upload new key** in the **Select Key** list and entered a user name and password in the **Administrator User Name** and **Administrator Password** fields.

**Important:** If you upload a new key for the existing Storwize V7000 user by using Tivoli Storage Productivity Center, that new key overrides the existing key for the user name. If you have another application that is using that user name (for example Tivoli Storage Productivity Center for Replication), the application cannot log on to Storwize V7000 because the key associated with the user is overwritten.

#### **Select User**

Select a user name from the list. The names are defined on Storwize V7000.

#### **Private SSH Key**

By default, the private key entry includes `${device.conf}\tpc_svc.pem`, which represents the Tivoli Storage Productivity Center default key `tpc_svc.pem`. The `tpc_svc.pem` key is in the `conf` directory where the Device server is installed. If you want to use your own key, the key must exist on the system where you are running the Tivoli Storage Productivity Center user interface. The key is uploaded to the Tivoli Storage Productivity Center server system. To identify a new key, click **Browse**.

The private SSH key must be in OpenSSH format or in PuTTY (.ppk) format that is not password protected. If you have a password protected PuTTY .ppk key, use the puttygen utility on a Windows operating system to convert that key to OpenSSH format by following these steps:

- a. Start the puttygen utility.
- b. Load the .ppk file. Click **File > Load private key**.
- c. From the menu bar, click **Conversion > Export OpenSSH key conversion**.

This action saves the key in OpenSSH format.

#### **Passphrase (Optional)**

Enter the passphrase for the private SSH key. The passphrase protects your private SSH key file. If you do not have a passphrase, leave this field blank.

### Storwize V7000 File Module (Optional)

If this check box *was not* selected when the storage system was added, you can select the check box to connect to the Storwize V7000 File Module. If this check box *was* selected when the storage system was added, you cannot clear the box.

Depending on the configuration of Storwize V7000 Unified, the information that is required to connect to this system can be the same as or different from the information for Storwize V7000. The fields and field descriptions are the same as the fields in the Storwize V7000 pane with the following exceptions:

- The **IP Address** field is not displayed because the IP address is retrieved automatically.
- The **User Name** field is always required.
- The private SSH key must be in OpenSSH format. The PuTTY format is not supported. You can use the puttygen utility to convert the key to the OpenSSH format.
- The **Select Key** list contains the option **Use user password** in addition to the **Upload new key** and **Use existing uploaded key** options. If you select **Use user password**, the **User Password** field is enabled and you can use a user name and password to connect to the file module.

If you select **Use user password**, the **Administrator User Name** and **Administrator Password** fields are optional.

3. Click **File > Save** on the menu bar. If the modifications are successful, the Storage Subsystems page is displayed and any new values for the storage subsystem are shown.

### Modifying System Storage DS8000 details

Learn how to modify information about a IBM System Storage DS8000 device that connects to IBM Tivoli Storage Productivity Center.

After you add a System Storage DS8000 device, basic information about that system appears on the **Administrative Services > Data Sources > Storage Subsystems** panel. To modify information about a System Storage DS8000 device on the Storage Subsystems panel, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources > Storage Subsystems**.
2. In the content panel, double-click a storage subsystem row or click the icon next to a row to edit information about that storage subsystem. Information about the storage subsystem appears on the Storage Subsystems Details window.
3. You can edit the information in the fields that are specified in bold in the window:

#### **Storage Subsystem**

Displays the name of the System Storage DS8000 device.

#### **HMC Address**

Displays the IP address or host name for the Hardware Management Console (HMC) that is used to manage the device.

#### **HMC2 Address (optional)**

Displays the IP address or host name for a second HMC that is used to manage the device.



**Username**

Displays the user name for the device.

**Password**

Displays the password for the device.

4. Select **File > Save**. If successful, the Storage Subsystems window appears and any new values for **Username** and **Password** for the storage subsystem are shown.

**Modifying other storage system details**

Learn how to modify information about the storage systems that connect to IBM Tivoli Storage Productivity Center using CIM object managers (CIMOMs).

After you add a CIMOM, basic information about that CIMOM appears on the **Administrative Services > Data Sources > Storage Subsystems** panel. To modify information about a CIMOM on the Storage Subsystems panel, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources > Storage Subsystems**.
2. In the content panel, double-click a storage subsystems row or click the icon next to a row to display a table of the CIMOMs that the storage subsystem is registered with.

**Service URL**

Displays the service Web address of the CIMOM. This Web address contains the IP address of the CIMOM, the port on which the CIMOM is listening, and the protocol used for communication. This Web address has a protocol (http | https), an IP address (IPv4 or IPv6), and a port number. For example: IPv4: https://127.0.0.1:5989 IPv6: https://2001:DB8:0:0:0:0:0:0:5989.

**Connection Status**

Displays the status of the connection between Tivoli Storage Productivity Center and the CIMOM.

**Interoperability Namespace**

Displays the interoperability namespace of the CIMOM. This namespace within the CIMOM allows the CIM Interop Schema (including the class instances of the Server Profile) to be accessed. The namespace also determines how Tivoli Storage Productivity Center interacts with the CIMOM when retrieving information. For a NetApp CIMOM, the value is /interop.

**Display Name**

Displays the name of the CIMOM as specified by the CIMOM provider.

**Description (optional)**

Displays the description that was entered on the Add CIMOM Agent window.

3. In the content panel, double-click a CIMOM row or click the icon next to a row to display details for the CIMOM. The connection information that you enter for a CIMOM is like the information that you enter on the Add CIMOM Agent panel. See “Adding a CIM agent connection” on page 244 for descriptions of the fields on this panel.
4. Select **File > Save**. If successful, the View and Edit CIMOM window appears and any new values for the CIMOM are shown.

## Testing the connection to a storage system

Learn how to test the connection from IBM Tivoli Storage Productivity Center to a storage system.

You must have Administrator privileges to perform this action.

After you add a storage system to Tivoli Storage Productivity Center, basic information about that system appears on the **Administrative Services > Data Sources > Storage Subsystems** panel. To verify that Tivoli Storage Productivity Center can use the information provided to connect to the storage system, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources > Storage Subsystems**.
2. In the content panel, highlight the row for a storage system and click **Test Connection** to test its connection to Tivoli Storage Productivity Center. The results of the test are displayed in the Test Subsystem Connection dialog box.
3. Click **OK** to close the dialog box. The **Connection Status** field is updated to reflect the status.

## Removing a storage system

Learn how to remove the connection between a storage system and IBM Tivoli Storage Productivity Center.

You must have Administrator privileges to perform this action.

Use the **Administrative Services > Data Sources > Storage Subsystems** panel to remove a storage system. Removing a storage system will:

- Remove the storage system information from master server database repository
- Remove the storage system information from all rollup reports
- Remove the storage system from the **Storage Subsystems** table

To remove a storage system, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources > Storage Subsystems**.
2. In the content panel, select a storage system and click **Remove**. A dialog box appears.
3. Click **Yes** to confirm the removal of the connection between the storage system and Tivoli Storage Productivity Center.

## VMware data source

VMware Virtual Infrastructure data sources are used by IBM Tivoli Storage Productivity Center to collect information from the hypervisors and virtual machines within them.

When you add a VMware data source in IBM Tivoli Storage Productivity Center, you need a user ID that has permission to browse through the data stores on VMware. IBM Tivoli Storage Productivity Center needs to browse through the data stores to collect information from the ESX Servers. However the "Read Only" role as defined by VMware does not allow IBM Tivoli Storage Productivity Center to browse the data stores. You can use the "Virtual Machine Power User" role as a choice if you do not want to use the Administrator role, or you can create a custom role with the required permissions.

To check what user role you have for VMware, follow these steps:

1. Connect a VI client to the ESX Server or VirtualCenter (whichever is used as the data source). Click the ESX Server and go to the **Summary** tab. Select a data source, right-click the data source and select **Browse**, to confirm that the browse permission is working correctly.
2. Connect the VI client to the data source with the admin role, go to the **Permission** tab and determine if the role name is used for the user.
3. Open the **Permissions** view, select that role, then right-click and select the **Edit** role. See what permissions are checked. Expand the check groups to find the specific ones.

For more information about VMware user roles, see: [http://www.vmware.com/pdf/vi3\\_vc\\_roles.pdf](http://www.vmware.com/pdf/vi3_vc_roles.pdf).

## Adding a VMware VI Data Source

This topic provides information about adding a VMware Virtual Infrastructure data source.

To add a VMware VI data source, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services > Data Sources > VMware VI Data Source**.
2. In the content pane, click **Add VMware VI Data Source**.
3. In the Add VMware VI Data Source window, enter the following information:

### Host Name

The host name for the VMware data source (ESX Server or VirtualCenter).

### Protocol

This is http or https. If you use https, you must have first imported the SSL certificate into IBM Tivoli Storage Productivity Center. For information about how to import the SSL certificate, see the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide Version 4 Release 1*. The protocol must match the same protocol that is set for the ESX Server or VirtualCenter.

**Port** The port used by the VMware data source.

### Username

The user name for the VMware data source.

### Password

The password for the VMware data source.

### Password Confirm

The password is confirmed for the VMware data source.

### Display Name

The name displayed in reports for the VMware data source.

### Description

The description for the VMware data source.

### Test VMware VI Data Source connectivity before adding

Check this box to check the connectivity for the VMware data source.

4. Click **Save**.

## Updating VMware data source information

After a VMware data source has been added, you can update the administrator name or password for the VMware data source.

To update information for a VMware data source, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources > VMware VI Data Source**.
2. In the content pane, click the icon to the left of the VMware data source you wish to change.
3. In the content pane, a notebook window opens. Make the appropriate changes.

## Displaying information about a VMware data source

You can view information about a VMware data source including its host name, protocol, port, and type.

To display information about a VMware data source, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources > VMware VI Data Source**.
2. In the content pane, click the icon to the left of the VMware data source for which you want information for. In the content pane, a notebook window opens. The following information is displayed:

### Host Name

The IP address of the data source.

### Protocol

This is http or https.

**Port** The port used by the VMware data source.

**Type** Virtual Center.

### Administrator Name

The name of the administrator for the VMware data source.

### Password

The password for the VMware data source.

### Password Confirm

Confirm the password for the VMware data source.

### Display Name

The name displayed for the VMware data source.

### Description

Description for the VMware data source.

### Software Level

The software level for the VMware data source.

### ConnectionStatus

The connection status to the VMware data source.

### Status Timestamp

### Test VMware VI Data Source connectivity before updating

Check this box to check the connectivity to the VMware data source.

## Show managed devices for VMware

You can view information about the managed devices for VMware.

To show the managed devices for VMware, follow this procedure:

1. In the navigation tree, expand **Administrative Services > Data Sources > VMware VI Data Source**.
2. In the content pane, click **Show Managed Devices**.
3. A pop-up window displays the VMware managed devices. Click **OK** to close the window.

## Deleting a VMware VI Data Source

This topic provides information about deleting a VMware Virtual Infrastructure data source.

To delete a VMware VI data source, follow this procedure:

1. In the Navigation Tree pane, expand **Administrative Services > Data Sources > VMware VI Data Source**.
2. In the content pane, click **Remove VMware VI Data Source**.
3. In the confirmation dialog window to remove a VMware data source, click **OK**.

---


## Managing schedules, jobs, and runs

Use the Job Management panel to view and manage all the schedules, runs, and jobs related to the storage devices that are monitored by IBM Tivoli Storage Productivity Center.

Specifically, you can use the Job Management panel to perform the following tasks:

- View information about all schedules, runs, and jobs in one central location. In previous versions of the product, this information was available only through nodes under each schedule type in the navigation tree.
- Identify all the schedules that include a particular storage device.
- Identify all the schedules, runs, and jobs that are encountering problems when executing.
- Ensure that a storage device is fully monitored by Tivoli Storage Productivity Center.
- Edit, rename, delete, and run schedules. You must be logged in as the user who created a schedule to perform some of these actions.

Use either of the following methods to access the Job Management panel:

- Click the job management icon  in the toolbar.
- Open the definition of a schedule and click the **View Job History** button to view the runs and jobs related to the schedule. For example, in the navigation tree, click **IBM Tivoli Storage Productivity Center > Monitoring > Probes > TPCUser.Default Probe** to view the probe definition panel and click **View Job History**.

The following terms are used on this panel:

- **Schedule:** A schedule is a definition of a monitoring task that collects data about storage resources. Monitoring tasks include scans, probes, pings, quotas, discovery, and performance monitors. For example, define a probe schedule to collect data about storage subsystems; define a scan schedule to collect file-level information about hosts.

- **Run:** A run is a single invocation of a schedule.
- **Job:** A job is a unit of work within a run for a schedule. One run might have multiple jobs. For example, a subsystem probe schedule has one job created per subsystem for which it collects data.

The Job Management panel is divided into the following sections:

#### **Filters section**

Use this section to perform the following tasks:

- Determine the type of storage entities whose schedules are displayed on the panel.
- View recommendations for storage entities that are not fully monitored by Tivoli Storage Productivity Center.

Filtering the contents of this panel based on entity type is optional. You select **Show All** from the **Entity Type** menu to display all the schedules defined in Tivoli Storage Productivity Center.

#### **Schedules section**

Use this section to manage the schedules related to the entities you selected in the **Filters** section. You can perform the following actions for each schedule:

- View and edit the definition of a schedule.
- Run a schedule immediately.
- Rename a schedule.
- Delete a schedule.
- Display schedules that have a certain status, such as Success, Running, or Failed.
- View basic information about a schedule, including the ID of the user who created it, description, job type (for example, ping, probe, or scan), the last time the schedule was run, and the next time the schedule is run.

#### **Jobs for Selected Schedule section**

Use this section to manage each run of a schedule and the jobs associated with a run. You can perform the following tasks for jobs and runs:

- View the status of a run and the date and time when it was executed.
- View information about the jobs associated with each run, including the name of the computer on which a job was run, the type of the monitored device, and the status of the job.
- View the log files for one or more jobs at a time.
- Display jobs that have started or completed within a specified length of time. For example, display jobs that have completed in the last two weeks.
- Display jobs that have a certain status, such as Success, Running, or Failed.

Click and drag the horizontal lines on the panel to resize each section. Click the names of columns in the **Schedules** and **Jobs for Selected Schedule** sections to sort the contents of lists based on the clicked column.

### **Filters**

#### **Refresh All**

Click this button to update all the lists that are displayed on the panel.

## Entity Type

Select the type of storage entity whose related schedule information you want to view. By default all schedules are displayed when you first access this panel. When you select a specific entity type, the following actions take place:

- The schedules, runs, and jobs related to the entity type are displayed in the appropriate sections on the panel.
- The names of the devices that belong to that entity type are displayed in the **Entity Name** list box.
- The text that is displayed above the **Show Recommendations** button indicates if the devices are fully monitored.

You can select from the following entity types in the **Entity Type** menu:

### Show All

(default) View all the schedules defined in Tivoli Storage Productivity Center.

### Storage Subsystem

View the schedules for storage subsystems that are monitored by Tivoli Storage Productivity Center.

**Fabric** View the schedules for fabrics that are monitored by Tivoli Storage Productivity Center.

### Computer

View the schedules for computers that are monitored by Tivoli Storage Productivity Center.

### Switch

View the schedules for switches that are monitored by Tivoli Storage Productivity Center.

### Tape Library

View the schedules for tape libraries that are monitored by Tivoli Storage Productivity Center.

### Monitoring Group

View the schedules that contain monitoring groups.

### Storage Resource Group

View the schedules for all the devices contained within a particular storage resource group.

## Entity Name Filter

Enter text to filter the names of entities that are displayed in the **Entity Name** list box. You can enter a combination of letters, numbers, dashes, underscores, and spaces. The filter text is not case sensitive. You can also enter the following wildcard characters:

- Use an asterisk (\*) to represent zero or more unknown characters of any value. For example, enter **DS8\*** to display devices of the selected type that begin with **DS** in their names, such as DS6000-A12-IBM, DS8000-B34-IBM, and DS5000-B35-IBM.
- Use a question mark to represent a single unknown character of any value. For example, enter **D?000-B3?-IBM** to include devices such as DS8000-B34-IBM, DS8000-B36-IBM, and DS5000-B35-IBM in the **Entity Name** list box.

### **Text that is displayed above the Show Recommendations button**

Warning text is displayed above the **Show Recommendations** button if the devices selected in the **Entity Name** list box are not fully monitored.

A device is considered not fully monitored if it is not included in all data collection schedules that are available in the licensed edition of Tivoli Storage Productivity Center that is installed. For example, if using the Standard edition of Tivoli Storage Productivity Center to monitor a host, that host must appear in the following schedule definitions to be considered fully monitored: ping, probe, and scan. If using the Disk edition of Tivoli Storage Productivity Center to monitor a storage subsystem, that storage subsystem must appear in the following schedule definitions to be considered fully monitored: probe.

If a device is considered not fully monitored, click **Show Recommendations** to view a list of recommended actions you can take to address the warning.

### **Show Recommendations**

Click this button to view recommended actions for a selected device that is not fully monitored. From the Recommendations panel, you can perform the following tasks:

- View the schedules that are recommended for a device.
- Click the **Take Action** button to create the recommended schedule for a device.

## **Schedules**

### **View Detail**

View the definition of a schedule that is selected in the schedule list. The schedule definition panel opens when you click this button. Use this panel to view and edit the definition of a schedule.

### **Run Now**

Run the selected schedule immediately. This option is only available if the job is not currently running.

### **Refresh**

Update the status of the schedule that is selected in the schedule list. The process to refresh the status of a schedule might take a few moments based on the current server load and network speed.

### **action drop down menu**

Use the options in this menu to perform actions against the selected schedule. The menu includes **View Detail**, **Run Now**, **Stop Now** (for performance monitor schedules) and the following actions:

#### **Rename**

Rename a schedule. You must be the owner of a schedule or the Tivoli Storage Productivity Center superuser to perform this action. This action is not available for the default schedules that are owned by TPCUser.

**Delete** Delete a schedule. When you delete a schedule, all the runs and jobs associated with the schedule are also deleted. You must be the owner of a schedule or the Tivoli Storage Productivity Center superuser to perform this action. This action is not available for the default schedules that are owned by TPCUser.



### Only show schedules with status

Filter the schedules that appear in the schedules list based on their status. For example, use this menu to indicate that only schedules with a status of **Running** are displayed in the schedules list. The default setting displays all schedules.

You can filter schedules using the following statuses:

- Show All
- Show All Unsuccessful
- Success
- Running
- Warnings
- Failed
- The job is waiting
- The job is running

Tivoli Storage Productivity Center automatically updates the schedules list if you change the value selected in this menu.

### schedules list

Use this list to view the schedules for the storage entities that are selected in the **Filters** section of the panel. The schedules that are displayed must also meet the filter selected in the **Only show schedules with status** menu. Click the name of a column to sort the contents of this list based on that column.

Some jobs that run in Tivoli Storage Productivity Center do not have a schedule associated with them. For example, jobs that are run through the command-line interface (CLI) do not have a schedule defined for them. To ensure that these job types are displayed on the Job Management panel, a new schedule called **CLI and Event Driven Jobs** was created. Select this schedule to view the runs and jobs that do not have user-created schedules associated with them. The job log files for the CLI and Event Driven Jobs schedule are automatically retained for 90 days and then deleted.

The following information is displayed for each schedule:

#### **Creator**

The ID of the Tivoli Storage Productivity Center user who created the schedule.

**Name** The name of a schedule.

#### **Description**

The description of a schedule.

#### **Job Type**

The type of the schedule, such as Probe, Ping, Scan, Aggregator Job, CIMOM Discovery, or Subsystem Performance Monitor.

#### **Last Run Finished**

The data and time when the most recent run of a schedule completed. This column is blank if the schedule has never been run.

#### **Last Run Status**

The status of the most recent run of a schedule.

#### **Next Run**

The date and time when the next run of the schedule is planned.

The value **On Demand** is displayed for schedules that are defined with the **Run Now** option. The value **Never** is displayed for schedules that are defined with the **Run Once** option and have already run.

**row number**

Appears in the lower right corner of the schedules list and indicates the number of rows in the schedules list. One row is displayed for each schedule. Use the horizontal lines on the panel to resize the section and increase or decrease the number of rows that are displayed.

## Jobs for Selected Schedule

**View Log File(s)**

View the log files of jobs that are selected in the runs list. Use the Shift + click and Ctrl + click keyboard shortcuts to select multiple jobs. You can view up to 10 logs at one time. The Log File Viewer panel opens when you click this button or when you double-click a row for a job. Use this panel to view all the processing performed by the Tivoli Storage Productivity Center server and agents for jobs.

**Note:** The number of job log files that are retained for a schedule is determined by the values defined on the Log File and Cached Batch Report Retention panel under **Administrative Services > Configuration**. However, for CLI and Event Driven Jobs schedules, job log files are automatically retained for 90 days.

**Show Jobs, in the last**

Use these fields to determine the runs that are displayed for the selected schedule. Use the following conditions in the filter:

- Select **Started** or **Completed** from the **Show Jobs** menu to display runs based on whether they started or completed in the time frame specified on the **last** field.
- Specify a time frame for started or completed runs by entering a number in the **last field** and selecting a unit of measurement from the menu.

For example, select **Started**, enter **4**, and select **weeks** to display the runs for a schedule that have started within the last four weeks. Tivoli Storage Productivity Center automatically updates the runs list if you change the value selected in these fields.

**Only show jobs with status**

Determine the runs that are displayed in the runs list based on their status. For example, use this menu to indicate that only runs with a status of **Completed with warning(s)** are displayed. The default setting displays all runs for the selected schedule.

You can filter the display of runs using the following statuses:

- Show All
- Show All Unsuccessful
- The job is waiting
- Running
- The job is running
- Completed successfully
- Running with warning(s)
- Completed with warning(s)

- Completed with error(s)
- System error occurred
- The agent did not respond (to a request from the server)
- The agent is being upgraded
- The agent is not licensed
- Agent not found
- Agent could not be reached
- The job delay limit was exceeded.
- The agent disconnected
- Host was shutdown or recycled
- Cluster resource group unreachable: no IP
- Scripts are disabled for this agent
- Script was not found or has an invalid extension
- The agent disk space is low

Tivoli Storage Productivity Center automatically updates the runs list if you change the value selected in this menu.

#### runs list

Use this list to view the runs and jobs of the schedule selected in the **Schedules** section of the panel. This list displays runs based on the criteria specified in the **Show jobs** and **in the last** fields.

Click the + symbol next to an entry in the **Run** field to view the jobs associated with a run. Click the name of a column to sort the contents of this list based on that column.

The following information is displayed for each run or job:

#### Run Status

The status of a run.

**Run** The run number assigned to a run and the date and time when the run was started.

#### Computer

(jobs only) The name of the computer where a job is executed. For database probes, this column displays the name of the Storage Resource agent that is performing the probe. For storage subsystem probes, this column displays **localhost** because the probe is performed by the Tivoli Storage Productivity Center server.

#### Device

(jobs only) The name of the device to which a job applies. This column is blank if the job is not related to a specific device.

**Status** (jobs only) The status of a job.

#### Start Time

(jobs only) The date and time when a job started.

#### Finish Time

(jobs only) The data and time when a job completed.

#### Log File Name

(jobs only) The name of the log file for a job. Click the **View Log File** to view the contents of a job log.

**row number**

Appears in the lower right corner of the runs list and indicates the number of rows in the runs list. One row is displayed for each run. You can expand each run row to view the jobs related to that run. Use the horizontal lines on the panel to resize the section and increase or decrease the number of rows that are displayed.

## Job Management - Recommendations

Use this window to view and apply the recommendations for storage devices that are not fully monitored by IBM Tivoli Storage Productivity Center.

A device is not fully monitored if it is not included in all the data collection schedules that are available in the licensed edition of Tivoli Storage Productivity Center that is installed. For example, if using the Standard Edition of Tivoli Storage Productivity Center to monitor a host, that host must appear in the following schedule definitions to be considered fully monitored: ping, probe, and scan. If using the Disk Edition of Tivoli Storage Productivity Center to monitor a storage subsystem, that storage subsystem must appear in a probe schedule to be considered fully monitored.

**Recommendation Severity**

Indicates whether a recommendation is a **Warning** or **Informational**.

- A **Warning** message indicates an action that you can take to more fully use Tivoli Storage Productivity Center to monitor a device. For example, when using a probe schedule in the Standard Edition of Tivoli Storage Productivity Center to collect data about a storage system, a warning message is displayed. This message recommends that you also include the storage subsystem in a performance monitor schedule. You can then use the data collected by performance monitor to identify the bottlenecks in storage usage and performance for the storage subsystem.
- An **Informational** message is a recommendation to upgrade Tivoli Storage Productivity Center to an edition that includes additional monitoring functions. For example, when using probe schedules in the Basic Edition of Tivoli Storage Productivity Center to collect data about a fabric, an informational message is displayed. This message indicates that if upgraded to the Standard Edition, Tivoli Storage Productivity Center can also collect performance data about the fabric.

For **Warning** messages, select the corresponding row and click **Take Action** to perform the recommended action.

**Entity** Displays the name of a storage entity to which the recommendation applies.

**Recommendation**

Indicates the recommended actions for a storage entity. One row is displayed per recommendation. Click **Take Action** to resolve recommendations that have a **Warning** severity.

If multiple recommendations are displayed for a device, sometimes performing one of the recommended actions might resolve the other recommendations. For example, consider if these two recommendations are displayed for a fabric:

- In order to collect information about device configuration, define a probe schedule for fabric {0}
- To enable zone control, add a data source that is capable of zone control for fabric {0}

If you click **Take Action** for the first recommendation and include the fabric in a probe, the second recommendation is also resolved. Both recommendations are then automatically removed from the **Recommendation** list

**Take Action**

Perform the recommended action for a warning. When you click this button, the corresponding panel for the action is displayed. For example, if the recommended action is to include a storage subsystem in a performance monitor, the panel is displayed.

**Cancel**

Exit the window without taking action.



---

## Chapter 4. Managing storage resources

The information in this chapter describes how to use IBM Tivoli Storage Productivity Center to manage your environment.

---

### Configure Devices

Use the Configure Devices wizard to set up storage devices for monitoring by IBM Tivoli Storage Productivity Center. The wizard guides you through the steps for adding a device as a data source, running a discovery, including devices in groups, specifying alerts, and setting up data collection schedules.

In previous versions of IBM Tivoli Storage Productivity Center, setting up a storage device for monitoring required performing tasks in different locations of the user interface, and each type of device required different steps. In Tivoli Storage Productivity Center v4.2 and later, all the tasks for setting up and configuring a device are included in the Configure Devices wizard.

Use the wizard to configure the following device types for monitoring:

- Storage subsystems
- Fabrics and switches
- Computers
- Tape libraries

You can add multiple devices of the same type using a single session of the wizard. For example, you can add an IBM DS8000, IBM XIV, IBM SAN Volume Controller, IBM Storwize V7000, and storage subsystems associated with CIM agents using a single session of the wizard. You cannot configure multiple devices of different types at the same time. For example, you cannot add a storage subsystem and a fabric using a single session of the wizard.

The pages in the wizard are divided into two sections:

#### Steps (left side)

Displays the steps for completing in the wizard. An icon is displayed next to the current step. The wizard includes the following steps for configuring a device:

1. Select the device type that you want to configure.
2. Select the specific devices that you want to configure.
3. Determine how data is collected from the devices.
4. Review your selections and run the wizard.
5. Review the results of the wizard.

How you access the wizard determines on what step you begin. For example, if you click the **Configure Devices** icon in the toolbar or select the **Configure Devices** node in the navigation tree, you start by selecting the device type that you want to configure. If you click the **Add Storage Subsystem** button on the Disk Manager > Storage Subsystem panel, you begin by entering connection information about the storage systems that you want to configure.

**Options (right side)**

Displays the options and fields that you must select and complete for each step.

**Related reference:**

“Alerts and schedules associated with monitoring groups and templates” on page 1137

Use alerts to notify you when conditions for storage devices are detected by a data collection schedule. Include storage devices in monitoring groups or templates to automatically associate alerts and schedules with those devices.

## Configure storage subsystems

Use the Configure Devices wizard to set up storage subsystems for monitoring by Tivoli Storage Productivity Center. Storage subsystems that you configure in the wizard are automatically included in probe schedules and associated with alert conditions.

You can configure one or more storage systems for monitoring by Tivoli Storage Productivity Center using a single session of the Configure Devices wizard. All the storage systems that you configure in a single pass through the wizard share the same probe schedule and alert conditions.

Depending on the type of storage system that you want to configure for monitoring, you must define connection information for its native interface or associated CIMOM. Enter information in the wizard about the native interfaces for the following storage systems:

- IBM System Storage DS8000 (native interface)
- IBM System Storage SAN Volume Controller (native interface)
- IBM XIV Storage System (native interface)
- IBM Storwize V7000 (native interface)
- IBM Storwize V7000 Unified (native interface)

For all other storage systems, you must enter information about the CIM agents that manage them.

To configure storage systems using the wizard, complete the following steps:

1. Access the wizard based on the task you want to perform:
  - If you want to add new or configure existing storage systems, click the configure devices icon in the toolbar or select the **Configure Devices** node in the navigation tree. You cannot configure existing storage systems if they are already included in a probe schedule. If you use this method, you must also complete the following steps:
    - a. Select **Storage Subsystem** on the Select Device Type page.
    - b. Determine whether to add new storage systems or configure existing storage systems that are known to Tivoli Storage Productivity Center.
      - To add storage systems, select **Add and configure new storage subsystems** and go to step 3.
      - To configure existing storage systems, click **Configure existing storage subsystems**. The page displays the storage systems that are known to Tivoli Storage Productivity Center. Go to step 2.
  - If you only want to add new storage systems, use any of the following methods and then go to step 3:



- In the navigation tree, go to **Disk Manager > Storage Subsystems** and click the **Add Storage Subsystem** button.
  - In the navigation tree, go to **IBM Tivoli Storage Productivity Center > Configuration Utility > Services > Storage Subsystem Interfaces** and click the **Add** button.
2. Select the existing storage systems that you want to configure. Storage systems that are already included in probe schedules are not displayed.  
If the Credentials column indicates **Missing** or **New Credentials Required** for the storage systems that you select, you must update their connection properties before continuing. See Configure storage subsystem connections page for descriptions of the connection properties.  
Go to step 6 after entering valid connection properties for each storage system that you want to configure.
  3. Configure storage system or CIMOM connections. Tivoli Storage Productivity Center uses the information you enter in this step to connect to a storage system or CIMOM. See Configure storage subsystem connections page for descriptions of the fields and options on the Configure storage system connections page.  
After entering information for a storage system or CIMOM, click **Add** to add it to the wizard session. Repeat the process to configure multiple storage systems or CIMOMs in the same session of the wizard.
  4. Discover the storage systems that are managed by the CIMOMs you entered in the previous step. Tivoli Storage Productivity Center collects initial configuration information for these entities.
  5. Select the storage systems that you want to configure. Any storage systems that were added in previous steps are automatically selected. Deselect storage systems that you do not want to configure.
  6. Specify how data is collected for storage systems by including them in a monitoring group or template. Use monitoring groups and templates to determine when data is collected and what alert conditions are checked for the storage systems that you are configuring.
    - Select a monitoring group if you want to manage a collection of storage systems in the same manner. By default, the monitoring group that you select is associated with an existing probe schedule and set of alerts. When you select the group, its data collection schedule and alerts are automatically applied to all the storage systems that you are configuring. After completing the wizard, the storage systems remain in the group and you can use that group when working in other parts of the user interface, such as in reports.

**Tip:** Tivoli Storage Productivity Center provides a number of default monitoring groups for storage systems. If you want to include a storage system in a user-defined group rather than in a default group, that group must be associated with a data collection schedule. If the user-defined group is not associated with a data collection schedule, it is not displayed in the list of monitoring groups from which you can select.

- Select a template if you want to create a unique probe schedule and set of alerts for a storage system rather than using the settings of an existing group.

The option that you select in this step determines the next step in the wizard.

- If you select a monitoring group, the following actions occur:
  - The storage systems are automatically included in a probe schedule to which the group is associated. For example, if you select Subsystem Basic

Group, the storage systems are included in the Subsystem Basic Probe schedule. The Subsystem Basic Probe schedule runs one time per week on Tuesday at 1:00 am.

- A set of alert conditions are applied to the storage systems in a group. When any of these conditions are detected during data collection, an alert is triggered. The group that you select determines the alerts that are associated with the storage systems. For example, if you select Subsystem Basic Group, 14 alert conditions are automatically checked for the storage systems.
- The storage systems are included in the group. You can use that group in other parts of the user interface, such as in reports.
- If you select a monitoring template, the following actions occur:
  - The storage systems are automatically included in a probe schedule that is created for the template. For example, if you select the Basic Subsystem Monitoring template, the storage systems are automatically included in a new probe that runs one time per week on Tuesday. The probe is created after the wizard is completed. Tivoli Storage Productivity Center determines the time when the probe is run based on server load to ensure optimal performance.
  - A set of alert conditions are applied to storage systems in the template. When any of these conditions are detected during data collection, an alert is triggered. The template you select determines the alerts that are associated with the storage systems. For example, if you select to the Basic Subsystem Monitoring template, 14 alert conditions are automatically checked for the storage systems.
  - You must enter a prefix for the probe schedule and alerts that are created for the template. This prefix is used to identify the probe and alerts in the user interface. For example, if you select the Basic Subsystem Monitoring template and enter the prefix "ABC", a probe schedule named "*userID*.ABCssBasicProbe" and alerts such as "*userID*.ABCPool Discovered (Basic)" are created, where *userID* represents the ID of the user that is logged in to Tivoli Storage Productivity Center.

The new probe and alerts are displayed under the corresponding node in the navigation tree for you to view and edit. For example, a probe schedule created from a template in the wizard appears under the **Tivoli Storage Productivity Center >Monitoring > Probes** node; new storage system alerts appear under the **Disk Manager > Alerting** node.

See Alerts and schedules associated with monitoring groups and templates for a list of groups and templates with their associated schedules and alerts.

7. Review the selections that you made in the wizard. The Review user selections page includes the following information:
  - A list of devices that you are configuring.
  - The name of the monitoring group that you selected. This value is not displayed if you selected a template.
  - The name of the probe schedule that is created based on the template you selected. This value is not displayed if you selected a monitoring group.
  - Information about the probe schedule created for a template.
  - The names of the alerts that are created based on the template you selected. This value is not displayed if you selected a monitoring group.
8. View the results of the device configuration and click **Finish** to close the wizard and probe the devices that you configured. The Review results page includes the following information:

- A list of the actions performed by the wizard. The page displays a row for the successful actions. If the configuration failed, a row is displayed for each failed action. For example, if the wizard expects to assign a specific alert to five devices, but the operation only succeeds for the three of the devices, this page displays one row for the successful actions and two rows each failed actions.
- Error messages for any failed actions.

## Configure storage subsystem connections page

Use this page in the Configure Devices wizard to enter or update connection properties for the storage systems and CIMOMs that you are configuring. Tivoli Storage Productivity Center uses this information to collect data about the storage systems.

Use the Configure storage system connections page to enter or update connection properties for the following storage entities. Depending on the type of storage system that you want to configure for monitoring, you must define connection information for the native interface or associated CIMOM for the storage system.

- System Storage DS8000 (native interface)
- SAN Volume Controller (native interface)
- XIV system (native interface)
- Storwize V7000 (native interface)
- Storwize V7000 Unified (native interface)
- Other (CIMOM)

The entity that is selected from the **Device Type** list determines the information that Tivoli Storage Productivity Center requires to connect to that entity. The following sections describe the fields for each entity.

**Requirement:** If you are adding or updating connection properties for a SAN Volume Controller or Storwize V7000 storage system, select **IBM SAN Volume Controller / IBM Storwize V7000** in the **Device Type** list. If you are adding or updating connection properties for a Storwize V7000 Unified storage system, select **IBM Storwize V7000 Unified**.

### IBM DS8000

Enter or update the following information to have Tivoli Storage Productivity Center connect to a System Storage DS8000 storage system.

#### HMC Address

Enter the IP address or host name for the Hardware Management Console (HMC) that is used to manage the System Storage DS8000 storage system.

#### HMC2 Address (Optional)

Enter the IP address or host name for a second HMC that is used to manage the storage system.

#### Username

Enter the user name for logging in to the System Storage DS8000 element manager. The default user name is admin.

**Note:** This user name is the same as the ESSNI user name.

#### Password

Enter the password for the user name.

## IBM XIV

Enter or update the following information to have Tivoli Storage Productivity Center connect to the XIV system.

### IP Address

Enter the IP address or host name for the XIV system that you want to configure.

### Username

Enter the user name for the XIV system.

### Password

Enter the password for the user name.

## IBM SAN Volume Controller / IBM Storwize V7000

Enter or update the following information to have Tivoli Storage Productivity Center connect to a SAN Volume Controller or Storwize V7000 storage system.

### Software Version

Enter the version of the SAN Volume Controller or Storwize V7000 software that you want to configure.

### IP Address

Enter the IP address or host name of the SAN Volume Controller or Storwize V7000. You can enter an IPv4 or IPv6 address, for example:

- 127.0.0.1
- [2001:DB8:0:0:0:0:0:0/32]

### Select Key

A Secure Shell (SSH) key is required to access the storage system. Select one of the following actions for the SSH key:

- **Upload new key:** Select this option to upload a new key. Public key content is extracted from the private key that is provided in the **Private SSH Key** field and is uploaded to the SAN Volume Controller or Storwize V7000 cluster. If you select this action, the **Admin Username** and **Admin Password** fields are required.
- **Use existing uploaded key:** Select this option if you want to use a public key that was already uploaded by using a method other than through Tivoli Storage Productivity Center, such as the SAN Volume Controller or Storwize V7000 web interface. Enter the location of the private key file that corresponds to the public key in the **Private SSH Key** field. The **Admin Username** and **Admin Password** fields are not required if you select this option.

### Admin Username

If you selected **Upload new key** in the **Select Key** list, enter the name of a user that has the Security Administrator role for the SAN Volume Controller or Storwize V7000 cluster that contains the storage system. Tivoli Storage Productivity Center uses this value during the setup process to contact the cluster and configure the SSH keys for the user. This user name must have privileges to modify other user accounts (Security Administrator role); otherwise, Tivoli Storage Productivity Center cannot configure the SSH keys.

If you selected **Use existing uploaded key**, this field is optional.

**Note:** The **Admin Username** only applies to SVC5.x+; it does not apply to SVC 4.x.

#### **Admin Password**

Enter the password for the user that you entered in the **Admin Username** field.

#### **Username**

Enter an authorized user for the storage system. Tivoli Storage Productivity Center uses this value when monitoring and configuring the storage system. If the user name does not exist it is created. This user name must have the Administrator role.

This field is active if you have selected **Upload new key** in the **Select Key** list and have entered a user name and password in the **Admin Username** and **Admin Password** fields.

**Important:** If you upload a new key for the existing SAN Volume Controller or Storwize V7000 user by using Tivoli Storage Productivity Center, that new key overrides the existing key for the user. If you have another application that is using that user (for example Tivoli Storage Productivity Center for Replication), the application cannot log on to SAN Volume Controller or Storwize V7000 because the key associated with the user is overwritten.

If you have this situation, create a user ID specifically for Tivoli Storage Productivity Center or Tivoli Storage Productivity Center for Replication on the SAN Volume Controller or Storwize V7000 cluster. Add the SAN Volume Controller or Storwize V7000 to Tivoli Storage Productivity Center for Replication using the new user ID and password (or change the user ID and password while the storage system is in the disconnected state).

#### **Select User**

Click this button to select from a list of existing users that are defined on the SAN Volume Controller or Storwize V7000.

#### **Private SSH Key**

By default, the private key entry includes `${device.conf}\tpc_svc.pem`, which represents the Tivoli Storage Productivity Center default key `tpc_svc.pem`. The `tpc_svc.pem` key is in the `conf` directory where the Device server is installed. If you want to use your own key, the key must exist on the system where you are running the Tivoli Storage Productivity Center user interface. The key is uploaded to the Tivoli Storage Productivity Center server system. To identify a new key, click **Browse**.

The private SSH key must be in OpenSSH format or in PuTTY (.ppk) format that is not password protected. If you have a password protected PuTTY .ppk key, use the `puttygen` utility on a Windows operating system to convert that key to OpenSSH format by following these steps:

1. Start the `puttygen` utility.
2. Load the .ppk file. Click **File > Load private key**.
3. From the menu bar, click **Conversion > Export OpenSSH key conversion**.

This action saves the key in OpenSSH format.

#### **Passphrase (Optional)**

Enter the passphrase for the private SSH key. The passphrase protects your private SSH key file. If you do not have a passphrase, leave this field blank.

## Storwize V7000 Unified

Enter or update the following information to have Tivoli Storage Productivity Center connect to a Storwize V7000 Unified storage system.

Storwize V7000 Unified is composed of Storwize V7000 and Storwize V7000 File Module. You can connect to only Storwize V7000 or to both Storwize V7000 and Storwize V7000 File Module. Storwize V7000 provides block-level data and Storwize V7000 File Module provides file-level data.

### Storwize V7000

Use the information in this section to connect Storwize V7000.

#### IP Address

Enter the IP address or host name of the Storwize V7000 storage system. You can enter an IPv4 or IPv6 address, for example:

- 127.0.0.1
- [2001:DB8:0:0:0:0:0:0/32]

#### Select Key

An SSH key is required to access Storwize V7000. Select one of the following actions for the SSH key:

- **Upload new key:** Select this option to upload a new key. Public key content is extracted from the private key that is provided in the **Private SSH Key** field and is uploaded to the Storwize V7000 cluster. If you select this action, the **Administrator User Name** and **Administrator Password** fields are required.
- **Use existing uploaded key:** Select this option if you want to use a public key that was already uploaded by using a method other than through Tivoli Storage Productivity Center, such as the Storwize V7000 web interface. Enter the location of the private key file that corresponds to the public key in the **Private SSH Key** field. The **Administrator User Name** and **Administrator Password** fields are not required if you select this option.

#### Administrator User Name

If you selected **Upload new key** in the **Select Key** list, enter the name of a user that has the Security Administrator role for the Storwize V7000 cluster that contains the storage system. Tivoli Storage Productivity Center uses this value during the setup process to contact the cluster and configure the SSH keys for the user. This user name must belong to the Security Administrator role, which has privileges to modify other user accounts. Otherwise, Tivoli Storage Productivity Center cannot configure the SSH keys.

If you selected **Use existing uploaded key**, this field is optional.

#### Administrator Password

Enter the password for the user name that you entered in the **Administrator User Name** field.

#### User Name

Enter an authorized user name for the storage system. Tivoli Storage Productivity Center uses this value to monitor and configure the storage system. If the user name does not exist, it is created. This user name must belong to the Administrator role.

This field is active if you selected **Upload new key** in the **Select Key** list and entered a user name and password in the **Administrator User Name** and **Administrator Password** fields. **Important:** If you upload a new key for the existing Storwize V7000 user by using Tivoli Storage Productivity Center, that new key overrides the existing key for the user name. If you have another application that is using that user name (for example Tivoli Storage Productivity Center for Replication), the application cannot log on to Storwize V7000 because the key associated with the user is overwritten.

#### Select User

Select a user name from the list. The names are defined on Storwize V7000.

#### Private SSH Key

By default, the private key entry includes `${device.conf}\tpc_svc.pem`, which represents the Tivoli Storage Productivity Center default key `tpc_svc.pem`. The `tpc_svc.pem` key is in the `conf` directory where the Device server is installed. If you want to use your own key, the key must exist on the system where you are running the Tivoli Storage Productivity Center user interface. The key is uploaded to the Tivoli Storage Productivity Center server system. To identify a new key, click **Browse**.

The private SSH key must be in OpenSSH format or in PuTTY (.ppk) format that is not password protected. If you have a password protected PuTTY .ppk key, use the puttygen utility on a Windows operating system to convert that key to OpenSSH format by following these steps:

1. Start the puttygen utility.
2. Load the .ppk file. Click **File > Load private key**.
3. From the menu bar, click **Conversion > Export OpenSSH key conversion**.

This action saves the key in OpenSSH format.

#### Passphrase (Optional)

Enter the passphrase for the private SSH key. The passphrase protects your private SSH key file. If you do not have a passphrase, leave this field blank.

#### Storwize V7000 File Module (Optional)

To connect to the Storwize V7000 Unified system, select this check box.

Depending on the configuration of Storwize V7000 Unified, the information that is required to connect to this system can be the same as or different from the information for Storwize V7000. The fields and field descriptions are the same as the fields in the Storwize V7000 pane with the following exceptions:

- The **IP Address** field is not displayed because the IP address is retrieved automatically.
- The **User Name** field is always required.
- The private SSH key must be in OpenSSH format. The PuTTY format is not supported. You can use the puttygen utility to convert the key to the OpenSSH format.
- The **Select Key** list contains the option **Use user password** in addition to the **Upload new key** and **Use existing uploaded key** options. If you

select **Use user password**, the **User Password** field is enabled and you can use a user name and password to connect to the file module.

If you select **Use user password**, the **Administrator User Name** and **Administrator Password** fields are optional.

## Other (CIMOM)

Enter or update connection properties for a CIMOM. Tivoli Storage Productivity Center uses this information to discover the storage systems that are managed by a CIMOM. You can then configure those storage systems as part of the wizard session.

The connection information that you enter for a CIMOM is similar to the information that you enter when adding a CIMOM on the Add CIMOM page. See "Adding a CIM agent connection" on page 244 for descriptions of the fields on this page.

## Discovered Storage Systems

After you run a discovery job, a table is displayed at the bottom of the "Configure storage subsystem connections" page. This table provides information about the storage systems that were discovered for the selected storage system type.

You can remove the connection to a storage system by clicking the row for the system in the table, and then clicking **Remove**.

## Configure fabrics and switches

Use the Configure Devices wizard to set up fabrics and switches for monitoring by Tivoli Storage Productivity Center. Fabrics and switches that you configure in the wizard are automatically included in probe schedules and associated with alert conditions.

You can use the wizard to configure the following fabric and switch types for monitoring:

- Brocade
- McData
- Cisco
- QLogic
- Brocade/McData (heterogeneous fabric)

You can configure one or more fabrics using a single session of the Configure Devices wizard. All the fabrics and switches that you configure in a single pass through the wizard share the same probe schedule and alert conditions.

To configure fabrics and switches using the wizard, complete the following steps:

1. Access the wizard based on the task you want to perform:
  - If you want to add new or configure existing fabrics and switches, click the configure devices icon in the toolbar or select the **Configure Devices** node in the navigation tree. If you use this method, you must also complete the following steps:
    - a. Select **Fabrics/Switches** on the Select Device Type page.
    - b. Determine whether to add new fabrics or configure existing fabrics that are known to Tivoli Storage Productivity Center.



- To add fabrics and switches, select **Add and configure new fabrics** and go to step 3.
  - To configure existing fabrics and switches, select **Configure existing fabrics**. The page displays the fabrics that are known to Tivoli Storage Productivity Center. Go to step 2.
- If you only want to add new fabrics and switches, go to **Fabric Manager > Fabrics** and click the **Add Fabric** button. Go to step 3.
2. Select the existing fabrics that you want to configure. Fabrics that are already included in probe schedules are not displayed. Go to step 6 after selecting the fabrics that you want to configure.
  3. Configure CIMOM or agent connections. The fields for entering connection information for a fabric changes based on the type of fabric you are adding. Tivoli Storage Productivity Center uses the information you enter in this step to connect to a fabric or CIMOM. Click **Add** to configure multiple fabrics in a session of the wizard.

**Note:** Depending on the type of fabric you are adding, the wizard recommends whether to enter CIMOM or agent information. For example, if you are configuring a Brocade fabric, the recommended way to administer it is with a CIMOM. If a CIMOM is not available, enter connection information for an out of band fabric agent.

4. Discover the fabrics and fabrics managed by the CIMOMs that you entered in the previous step. Tivoli Storage Productivity Center collects initial configuration information for these entities.
5. Select the fabrics that you want to configure for monitoring. Any fabrics that were added in previous steps are automatically selected. Deselect fabrics that you do not want to configure.
6. Specify how data is collected for fabrics by including them in a monitoring group or template. Use monitoring groups and templates to determine when data is collected and what alert conditions are checked for the fabrics that you are configuring.
  - Select a monitoring group if you want to manage a collection of fabrics in the same manner. Each monitoring group that you can select is associated with an existing probe schedule and set of alerts. When you select a group, its data collection schedule and alerts are automatically applied to all the fabrics that you are configuring. After completing the wizard, the fabrics remain in the group and you can use that group when working in other parts of the user interface, such as in reports.

**Tip:** Tivoli Storage Productivity Center provides a number of default monitoring groups for fabrics. If you want to include a fabric in a user-defined group rather than in a default group, that group must be associated with a data collection schedule. If the user-defined group is not associated with a data collection schedule, it is not displayed in the list of monitoring groups from which you can select.

- Select a template if you want to create a unique probe schedule and set of alerts for a fabric rather than using the settings of an existing group.

The option that you select in this step determines the field and pages that appear next in the wizard.

- If you select a monitoring group, the following actions occur:
  - The fabrics are automatically included in a probe schedule to which the group is associated. For example, if you select Fabric Basic Group, the

fabrics are included in the Fabric Basic Probe schedule. The Fabric Basic Probe schedule runs two times per week on Monday and Wednesday at 3:00 am.

- A set of alert conditions are applied to the fabrics in a group. When any of these conditions are detected during data collection, an alert is triggered. The group that you select determines the alerts that are associated with the fabrics. For example, if you select Fabric Basic Group, nine alert conditions are automatically checked for the fabrics.
- The fabrics are included in the group. You can use that group in other parts of the user interface, such as in reports.
- If you select a monitoring template, the following actions occur:
  - The fabrics are automatically included in a probe schedule that is created for a template. For example, if you select the Basic Fabric Monitoring template, the fabrics are automatically included in a new probe that runs two times per week on Monday and Wednesday. The probe is created after the wizard is completed. Tivoli Storage Productivity Center determines the time when the probe is run based on server load to ensure optimal performance.
  - A set of alert conditions are applied to the fabrics in a template. When any of these conditions are detected during data collection, an alert is triggered. The template you select determines the alerts that are associated with the fabrics. For example, if you select to the Basic Fabric Monitoring template, nine alert conditions are automatically checked for the fabrics.
  - You must enter a prefix for the probe schedule and alerts that are based on a template. This prefix is used to identify the probe and alerts in the user interface. For example, if you select the Basic Fabric Monitoring template and enter the prefix "ABC", a probe schedule named "*userID*.ABCfabBasicProbe" and alerts such as "*userID*.ABCFabric State Change (Basic)" are created, where *userID* represents the ID of the user that is logged in to Tivoli Storage Productivity Center.

The new probe and alerts are displayed under the corresponding node in the navigation tree for you to view and edit. For example, a probe schedule created from a template in the wizard appears under the **Tivoli Storage Productivity Center >Monitoring > Probes** node; new storage system alerts appear under the **Fabric Manager > Alerting** node.

See Alerts and schedules associated with monitoring groups and templates for a list of groups and templates with their associated schedules and alerts.

7. Review the selections that you made in the wizard. The Review user selections page includes the following information:
  - A list of devices that you are configuring.
  - The name of the monitoring group that you selected. This value is not displayed if you selected a template.
  - The name of the probe schedule that is created based on the template you selected. This value is not displayed if you selected a monitoring group.
  - Information about the probe schedule created for a template.
  - The names of the alerts that are created based on the template you selected. This value is not displayed if you selected a monitoring group.
8. View the results of the device configuration and click **Finish** to close the wizard and probe the devices that you configured. The Review results page includes the following information:
  - A list of the actions performed by the wizard. The page displays a row for the successful actions. If the configuration failed, a row is displayed for each

failed action. For example, if the wizard expects to assign a specific alert to five devices, but the operation only succeeds for the three of the devices, this page displays one row for the successful actions and two rows each failed actions.

- Error messages for any failed actions.

## Configure computers

Use the Configure Devices wizard to set up computers for monitoring by Tivoli Storage Productivity Center. Computers that you configure in the wizard are automatically included in probe and ping schedules and associated with alert conditions.

Use the Configure Devices wizard to configure computers that meet the following conditions:

- A Storage Resource agent has been deployed on the computer.
- The computer is not included in an existing probe schedule.

You can configure one or more computers using a single session of the Configure Devices wizard. All the computers that you configure in a single pass through the wizard share the same probe schedule, ping schedule, and alert conditions.

To configure computers using the wizard, complete the following steps:

1. Click the configure devices icon in the toolbar or select the **Configure Devices** node in the navigation tree.
2. Select **Computers** on the Select Device Type page.
3. Select the computers that you want to configure. Computers that are already included in an existing probe schedule are not displayed.
4. Specify how data is collected for computers by including them in a monitoring group or template. Use monitoring groups and templates to determine when data is collected and what alert conditions are checked for the computers that you are configuring.
  - Select a monitoring group if you want to manage a collection of computers in the same manner. Each monitoring group that you can select is associated with existing probe and ping schedules and a set of alerts. When you select a group, its data collection schedules and alerts are automatically applied to all the computers that you are configuring. After completing the wizard, the computers remain in the group and you can use that group when working in other parts of the user interface, such as in reports.

**Tip:** Tivoli Storage Productivity Center provides a number of default monitoring groups for computers. If you want to include a computer in a user-defined group rather than in a default group, that group must be associated with a data collection schedule. If the user-defined group is not associated with a data collection schedule, it is not displayed in the list of monitoring groups from which you can select.

- Select a template if you want to create a unique probe schedule, ping schedule, and set of alerts for a computer rather than using the settings of an existing group.

The option that you select in this step determines the field and pages that appear next in the wizard.

- If you select a monitoring group, the following actions occur:

- The computers are automatically included in probe and ping schedules to which the group is associated. For example, if you select Computer Basic Group, the computers are included in the Computer Basic Probe and Computer Basic Ping schedules. The Computer Basic Probe schedule runs two times per week on Monday and Wednesday at 3:00 am. The Computer Basic Ping schedule runs every 3 hours.
- A set of alert conditions are applied to the computers in a group. When any of these conditions are detected during data collection, an alert is triggered. The group that you select determines the alerts that are associated with the computers. For example, if you select Computer Basic Group, seven alert conditions are automatically checked for the computers.
- The computers are included in the group. You can use that group in other parts of the user interface, such as in reports.
- If you select a monitoring template, the following actions occur:
  - The computers are automatically included in a probe and ping schedule that are created for the template. For example, if you select the Basic Computer Monitoring template, the computers are automatically included in a new probe that runs two times per week on Monday and Wednesday and a ping that runs every three hours. The probe and ping are created after the wizard is completed. Tivoli Storage Productivity Center determines the time when the probe is run based on server load to ensure optimal performance.
  - A set of alert conditions are applied to the computers in a template. When any of these conditions are detected during data collection, an alert is triggered. The template you select determines the alerts that are associated with the computers. For example, if you select to the Basic Computer Monitoring template, seven alert conditions are automatically checked for the computers.
  - You must enter a prefix for the probe schedules, ping schedules, and alerts that are based on a template. This prefix is used to identify these definitions in the user interface. For example, if you select the Standard Computer Monitoring template and enter the prefix "ABC", the following definitions are created in the user interface after the wizard is completed:
    - a probe schedule named "*userID*.ABCcompStandardProbe"
    - a ping schedule named "*userID*.ABCcompStandardPing"
    - seven alerts that follow this naming convention: *userID*.ABCStatus Change Online (Standard)

Where *userID* represents the ID of the user that is logged in to Tivoli Storage Productivity Center.

The new schedule definitions and alerts are displayed under the corresponding node in the navigation tree for you to view and edit. For example, a probe schedule created from a template in the wizard is displayed under the **Tivoli Storage Productivity Center >Monitoring > Probes** node; a ping schedule is displayed under the **Data Manager > Monitoring > Pings** node; new storage system alerts appear under the **Data Manager > Alerting** node.

See Alerts and schedules associated with monitoring groups and templates for a list of groups and templates with their associated schedules and alerts.

5. Review the selections that you made in the wizard. The Review user selections page includes the following information:
  - A list of devices that you are configuring.

- The name of the monitoring group that you selected. This value is not displayed if you selected a template.
  - The name of the probe and ping schedules that are created based on the template you selected. This value is not displayed if you selected a monitoring group.
  - Information about the schedules created for a template.
  - The names of the alerts that are created based on the template you selected. This value is not displayed if you selected a monitoring group.
6. View the results of the device configuration and click **Finish** to close the wizard and probe the devices that you configured. The Review results page includes the following information:
- A list of the actions performed by the wizard. The page displays a row for the successful actions. If the configuration failed, a row is displayed for each failed action. For example, if the wizard expects to assign a specific alert to five devices, but the operation only succeeds for the three of the devices, this page displays one row for the successful actions and two rows each failed actions.
  - Error messages for any failed actions.

## Configure tape libraries

Use the Configure Devices wizard to set up tape libraries for monitoring by IBM Tivoli Storage Productivity Center. Tape libraries that you configure in the wizard are automatically included in probe schedules.

You can configure one or more tape libraries using a single session of the Configure Devices wizard. All the tape libraries that you configure in a single pass through the wizard share the same probe schedule.

To configure tape libraries using the wizard, complete the following steps:

1. Access the wizard based on the task you want to perform:
  - If you want to add new or configure existing tape libraries, click the configure devices icon in the toolbar or select the **Configure Devices** node in the navigation tree. If you use this method, you must also complete the following steps:
    - a. Select **Tape Libraries** on the Select Device Type page.
    - b. Determine whether to add new tape libraries or configure existing tape libraries that are known to Tivoli Storage Productivity Center.
      - To add tape libraries, select **Add and configure new tape libraries** and go to step 3.
      - To configure existing tape libraries, select **Configure existing tape libraries**. The page displays the tape libraries that are known to Tivoli Storage Productivity Center. Go to step 2.
  - If you only want to add new tape libraries, go to **Tape Manager > Tape Libraries** and click the **Add Tape Library** button. Go to step 3.
2. Select the existing tape libraries that you want to configure. Tape libraries that are already included in probe schedules are not displayed. Go to step 6 after selecting the tape libraries that you want to configure.
3. Configure CIMOM connections for tape libraries. Tivoli Storage Productivity Center uses the information you enter in this step to connect to a CIMOM. Click **Add** to configure multiple CIMOMs in a session of the wizard.

The connection information that you enter for a CIMOM is like the information that you enter on the Add CIMOM panel.

4. Discover the tape libraries that are managed by the CIMOMs you entered in the previous step. Tivoli Storage Productivity Center uses the information you enter in this step to connect to a CIMOM. Click **Add** to configure multiple CIMOMs in a session of the wizard.
5. Select the tape libraries that you want to configure for monitoring. Any tape libraries that were added in previous steps are automatically selected. Deselect tape libraries that you do not want to configure.
6. Specify how data is collected for tape libraries by including them in a monitoring group or template. Use monitoring groups and templates to determine when data is collected for the tape libraries that you are configuring.
  - Select a monitoring group if you want to manage a collection of tape libraries in the same manner. Each monitoring group that you can select is associated with an existing probe schedule. When you select a group, its data collection schedule is automatically applied to all the tape libraries that you are configuring. After completing the wizard, the tape libraries remain in the group and you can use that group when working in other parts of the user interface, such as in reports.

**Tip:** Tivoli Storage Productivity Center provides a number of default monitoring groups for tape libraries. If you want to include a tape library in a user-defined group rather than in a default group, that group must be associated with a data collection schedule. If the user-defined group is not associated with a data collection schedule, it is not displayed in the list of monitoring groups from which you can select.

- Select a template if you want to create a unique probe schedule for a tape library rather than using the settings of an existing group.

The option that you select in this step determines the field and pages that appear next in the wizard.

- If you select a monitoring group, the following actions occur:
  - The tape libraries are automatically included in a probe schedule to which the group is associated. For example, if you select Tape Basic Group, the tape libraries are included in the Tape Basic Probe schedule. The Tape Basic Probe schedule runs one time per week on Sunday at 4:00 am.
  - The tape libraries are included in the group. You can use that group in other parts of the user interface, such as in reports.
- If you select a monitoring template, the following actions occur:
  - The tape libraries are automatically included in a probe schedule that is created for a template. For example, if you select the Basic Tape Monitoring template, the tape libraries are automatically included in a new probe that runs one time per week on Sunday. The probe is created after the wizard is completed. Tivoli Storage Productivity Center determines the time when the probe is run based on server load to ensure optimal performance.
  - You must enter a prefix for the probe schedule and alerts that are based on a template. This prefix is used to identify the probe and alerts in the user interface. For example, if you select the Basic Tape Monitoring template and enter the prefix "ABC", a probe schedule named "*userID*.ABCtapeBasicProbe" is created, where *userID* represents the ID of the user that is logged in to Tivoli Storage Productivity Center.

The new probe is displayed under the corresponding node in the navigation tree for you to view and edit. For example, a probe schedule created from a template in the wizard appears under the **Tivoli Storage Productivity Center >Monitoring > Probes** node.

See Alerts and schedules associated with monitoring groups and templates for a list of groups and templates with their associated schedules and alerts.

7. Review the selections that you made in the wizard. The Review user selections page includes the following information:
  - A list of devices that you are configuring.
  - The name of the monitoring group that you selected. This value is not displayed if you selected a template.
  - The name of the probe schedule that is created based on the template you selected. This value is not displayed if you selected a monitoring group.
  - Information about the probe schedule created for a template.
  - The names of the alerts that are created based on the template you selected. This value is not displayed if you selected a monitoring group.
8. View the results of the device configuration and click **Finish** to close the wizard and probe the devices that you configured. The Review results page includes the following information:
  - A list of the actions performed by the wizard. The page displays a row for the successful actions. If the configuration failed, a row is displayed for each failed action. For example, if the wizard expects to assign a specific alert to five devices, but the operation only succeeds for the three of the devices, this page displays one row for the successful actions and two rows each failed actions.
  - Error messages for any failed actions.

---

## Discovering storage resources

This topic describes how you can collect details about storage resources in your environment.

The process of finding storage resources within an enterprise, including detection of network topology, is called discovery. Use discovery schedules in Tivoli Storage Productivity Center to collect basic information about the storage resources in your environment, including computers, CIMOMs, storage systems, fabrics and switches, NetWare trees, filers, clusters, and tape libraries. You must run discovery schedules before you can collect more detailed information about storage resources using other data collection schedules, such as probes, scans, and pings.

You are not required to run discovery schedules for storage resources that meet any of the following conditions:

- Storage resources that connect to Tivoli Storage Productivity Center using native interfaces.
- Storage resources that have been configured in the Configure Devices wizard.

## Discovery overview

IBM Tivoli Storage Productivity Center supports various discovery types. Use discovery schedules to find new CIMOMs, Out of Band Fabrics, NetWare Filers, Windows Domains, NAS filers, and SAN File System that have been introduced into your environment.

The discovery types in Tivoli Storage Productivity Center are summarized in the following table. Additionally the table summarizes the types of information that each discovery provides.

*Table 16. Discovery types, their data sources, and top-level entities that are discovered*

| Type of discovery                                | Data source                                     | Top-level entities                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------------------------|-------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CIMOM                                            | CIM Agent                                       | <ul style="list-style-type: none"> <li>• Fabric (switches)</li> <li>• Storage subsystems</li> <li>• Tape subsystems</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Out-of-band                                      | SNMP agents on the switches                     | Fabric (switches). A fabric discovery gets all the fabric information switches that have been added as out-of-band agents. All information is collected by using the SNMP protocol to send queries across the IP network to the management information bases (MIBs) supported on the switch.                                                                                                                                                                                                                                                                                                                                         |
| Storage Resource agents or In-band Fabric agents | Storage Resource agent or In-band Fabric agents | <p>Fabric (switches). If you have deployed Storage Resource agents or have In-band Fabric agents on the computers attached to the fabric, you do not need to perform a specific discovery step for these Storage Resource agent. The fabric is automatically discovered.</p> <p><b>Note:</b> In Tivoli Storage Productivity Center v4.2 and above, you can deploy Storage Resource agents only. If you want to install an Inband fabric agent (Fabric agent), you must own a previous version of the product. See the documentation for the previous version of the product for information about how to install a Fabric agent.</p> |
| NetWare Filer                                    | NetWare agent                                   | NetWare trees                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Windows domain, NAS, and SAN File System         | Data agent or Storage Resource agent            | <ul style="list-style-type: none"> <li>• Computers</li> <li>• Filers</li> <li>• Clusters</li> <li>• Names or agents</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| VMWare VI                                        | VMWare VI data source                           | The data source can be a hypervisor (ESX Server or VirtualCenter).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

Generally, discovery only provides information about the existence of a resource. Use probes and scans to obtain detailed information about a specific resource after it has been discovered.



For a Windows domain, Tivoli Storage Productivity Center only performs a discovery on the first Windows system on the domain; not all the systems on the domain.

### **CIMOM discovery**

This topic describes how the IBM Tivoli Storage Productivity Center manages and monitors disk, switch, and tape devices through Common Information Model Object Managers (CIMOMs).

**Note:** CIMOM can be referred to by a variety of names, such as CIM agents or SMIS agents. A CIM agent consists of a CIMOM and an SMIS provider for the managed device.

CIM agents on the same subnet as the Tivoli Storage Productivity Center Device server are automatically discovered with CIMOM discovery. If CIM agents are not in the same subnet, they can be discovered using a Service Location Protocol (SLP) Directory Agent (DA) or they can be added manually. If you do not want Tivoli Storage Productivity Center to discover CIM agents in the local subnet, clear the Scan local subnet box in the CIMOM discovery Options panel. To include SLP Directory Agents in the discovery, add the IP address of the DA in the CIMOM discovery Options panel and click **Add**. CIMOM discovery has several purposes. It can scan the local subnet to find CIM agents and contact CIM agents that have been added through the dialog. After it has found a CIM agent, it attempts to log in and discover the storage entities being managed by the CIM agent. If a CIM agent is discovered but requires login information, you need to enter it in the View and Edit CIMOM panel and then rerun the discovery to collect the information for the storage entities. The amount of time CIMOM discovery takes depends on the number of CIM agents, the number of storage entities, and whether you are scanning the local subnet.

CIMOM discovery can be run on a schedule. How often you run it depends on how dynamic your environment is. It must be run to detect a new subsystem. CIMOM discovery also performs basic health checks of the CIM agent and subsystem.

For Tivoli Storage Productivity Center to successfully communicate with the CIM agents, the CIMOM service must be accessible through the IP network. The TCP/IP network configuration on the host where Tivoli Storage Productivity Center is installed must include in its list of domain names, all the domains that contain storage devices that are discoverable by Tivoli Storage Productivity Center.

The CIM agents are installed outside of Tivoli Storage Productivity Center. Contact your device vendor for information about installing the CIM agent.

#### **Limitation when running a CIMOM Discovery against a McDATA CIM agent:**

When you run a CIMOM discovery against a McDATA CIM agent, the following switch information is collected:

- Manageable McDATA switches that exist in the fabric.
- Unmanageable switches that connect to the manageable McDATA switches. In a homogenous fabric, a McDATA switch is considered unmanageable when it is managed by an Enterprise Fabric Connectivity Manager (Proxy Mode) or a McDATA CIM agent (Direct Mode). In a heterogenous fabric, an unmanageable switch is either a non-McDATA switch or a McDATA switch which is managed by different Enterprise Fabric Connectivity Manager in proxy mode or a different McDATA CIM agent in direct mode

You can use Tivoli Storage Productivity Center to view detailed information about the manageable McDATA switches and collect performance data about them. For unmanageable switches, you can use Tivoli Storage Productivity Center to view their WWN and port numbers which connect to manageable McDATA switches only, and you cannot collect performance data about them.

### SLP registration with slptool

This topic describes how IBM Tivoli Storage Productivity Center uses Service Location Protocol (SLP) discovery.

SLP discovery requires that all of the CIMOMs that Disk Manager discovers are registered using the SLP. However, for CIMOMs outside of the IP subnet, you need to use an SLP DA and register the CIMOM using **slptool**. In a non-multicast network environment, you must use the **slptool** as SLP can only discover CIMOMs that are registered in its IP subnet. When you use **slptool**, ensure that the **CIM\_InteropSchemaNamespace** and **Namespace** attributes are specified.

For example, enter the following command:

```
slptool register service:wbem:https://<myhost.com>:<port>
```

Where *myhost.com* is the name of the machine hosting the CIMOM, and *port* is the port number of the service, for example 5989.

**Note:** **slptool** is installed with a CIMOM. Run the command from the machine hosting the CIMOM.

### SMI-S discovery

This topic describes how the IBM Tivoli Storage Productivity Center can discover storage devices that comply with the Storage Management Initiative Specification (SMI-S).

SMI-S-compliant CIM Agents are discovered using the Service Location Protocol (SLP). The Tivoli Storage Productivity Center server software performs SLP discovery on the network. The SLP User Agent (UA) looks for all registered services with a service type of **service:wbem**.

The Tivoli Storage Productivity Center performs the following discovery tasks:

- Locates individual storage devices
- Retrieves vital characteristics for those storage devices
- Populates the internal Tivoli Storage Productivity Center databases with the discovered information

### Device server and discovery on a SAN

Storage Management Initiative (SMI) agents, legacy in-band fabric agents, out-of-band Simple Network Management Protocol (SNMP) agents, and Storage Resource agents can be used to discover and probe switches and fabrics. When multiple agents are available, IBM Tivoli Storage Productivity Center creates a list of the agents. The first agent on the list is started first. If an agent fails, the next agent on the list is started.

Host and device information is gathered by in-band or Storage Resource agents. Topology information is gathered using SMI (Brocade and McData switches only), in-band, out-of-band, or Storage Resource agents. The topology discovery

operations of all agent types provide a similar level of information. However, zone information is available only through certain agent types for certain vendors. See Table 17 for details.

For Brocade and McDATA fabrics, the preferred type of agent is the SMI agent. The SMI agent provides most fabric functions and the other agent types can be added for redundancy. However, Storage Resource agents or in-band fabric agents are required to gather host bus adapter (HBA) information.

For QLogic and Cisco fabrics, a combination of different agent types is required to enable all functions. If Storage Resource agents or in-band fabric agents are not used, the devices and hosts are displayed as entities of unknown type. You can change the icon and labels of the unknown entities to something more appropriate for your enterprise.

If you want to use out-of-band fabric agents, you must configure the switches to send SNMP traps to the the Device server. Similar configuration is not required for CIM indications for SMI agents because they are usually automatically subscribed to after a CIMOM discovery. However, in some environments, manual configuration of certain settings might be required. SNMP traps and CIM indications determine when to start an automatic rediscovery operation

The following table shows the actions that trigger data collection and the data that is collected depending on the type of agent that is used.

Table 17. Information collected by agent

| Action that triggers data collection                     | In-band fabric agent or Storage Resource agent | Out-of-band fabric agent (SNMP)                                                                                                                                                                            | SMI agent (Brocade and McData)                                                                                                                                                                   | SMI agent (Cisco and QLogic)                                                                                                                                         |
|----------------------------------------------------------|------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Administrative Services > Discovery > Out of Band Fabric | –                                              | <ul style="list-style-type: none"> <li>• Discovers switches and connections within an IP range</li> <li>• Retrieves switch and topology information</li> <li>• Retrieves Cisco VSAN information</li> </ul> | –                                                                                                                                                                                                | –                                                                                                                                                                    |
| Administrative Services > Discovery > CIMOM              | –                                              | –                                                                                                                                                                                                          | <ul style="list-style-type: none"> <li>• Discovers SMI agents through Service Location Protocol (SLP)</li> <li>• Retrieves switch information</li> <li>• Retrieves fabric information</li> </ul> | <ul style="list-style-type: none"> <li>• Discovers SMI agents through SLP</li> <li>• Retrieves switch information</li> <li>• Retrieves fabric information</li> </ul> |

Table 17. Information collected by agent (continued)

| Action that triggers data collection                                                                                                                                                                  | In-band fabric agent or Storage Resource agent                                                                                                                                                                                                                                                  | Out-of-band fabric agent (SNMP)                                                                                                           | SMI agent (Brocade and McData)                                                                                                                                                                                                                                                                                       | SMI agent (Cisco and QLogic)                                                                                              |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>IBM Tivoli Storage Productivity Center &gt; Probe</b><br>(including fabric)                                                                                                                        | <ul style="list-style-type: none"> <li>• Retrieves zoning information, excluding zone alias information, for QLogic, Cisco, and McDATA</li> <li>• Retrieves switch and topology information</li> <li>• Retrieves attribute information for hosts and devices connected to the fabric</li> </ul> | <ul style="list-style-type: none"> <li>• Retrieves switch and topology information</li> <li>• Retrieves Cisco VSAN information</li> </ul> | <ul style="list-style-type: none"> <li>• Retrieves detailed fabric, switch, switch port, and blade information</li> <li>• Retrieves zoning information, including zone aliases for Brocade devices</li> <li>• Retrieves fabric topology connectivity information</li> <li>• Subscribes to CIM indications</li> </ul> | <ul style="list-style-type: none"> <li>• Retrieves detailed fabric, switch, switch port, and blade information</li> </ul> |
| Action of in-band fabric agent or Storage Resource agent <sup>1</sup> : <ul style="list-style-type: none"> <li>• Turns on</li> <li>• Connectivity changes</li> <li>• Change event detected</li> </ul> | <ul style="list-style-type: none"> <li>• Retrieves zoning information, excluding zone alias information</li> <li>• Retrieves switch and topology information</li> <li>• Retrieves attribute information for hosts and devices connected to the fabric</li> </ul>                                | <ul style="list-style-type: none"> <li>• Retrieves switch and topology information</li> <li>• Retrieves Cisco VSAN information</li> </ul> |                                                                                                                                                                                                                                                                                                                      |                                                                                                                           |
| SNMP trap received from switch                                                                                                                                                                        | <ul style="list-style-type: none"> <li>• Retrieves zoning information, excluding zone alias information</li> <li>• Retrieves switch and topology information</li> <li>• Retrieves attribute information for hosts and devices connected to the fabric</li> </ul>                                | <ul style="list-style-type: none"> <li>• Retrieves switch and topology information</li> <li>• Retrieves Cisco VSAN information</li> </ul> |                                                                                                                                                                                                                                                                                                                      |                                                                                                                           |

Table 17. Information collected by agent (continued)

| Action that triggers data collection                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | In-band fabric agent or Storage Resource agent | Out-of-band fabric agent (SNMP) | SMI agent (Brocade and McData)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | SMI agent (Cisco and QLogic) |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| CIM indication received from fabric SMI agent                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                                                |                                 | Performs a mini-probe to collect information that is relevant to the indication received, or sets entities as missing. Might do one of the following actions: <ul style="list-style-type: none"> <li>• Retrieves switches and fabrics, or</li> <li>• Retrieves zoning data, or</li> <li>• Retrieves switch port status and connection to node, or</li> <li>• Retrieves switch blade status and all port connections for the blade, or</li> <li>• Sets fabric, switch, blades, connections, or nodes as missing</li> </ul> |                              |
| <p>1. The Storage Resource agent has two distinct functions:</p> <ul style="list-style-type: none"> <li>• <b>Discovery</b><br/>When a discovery job is run, Tivoli Storage Productivity Center gathers information about the fabrics, fabric to switch relationships, and other key attributes about the switches in the fabrics.</li> <li>• <b>Probe</b><br/>When a probe job is run, Tivoli Storage Productivity Center gathers additional details for topology and zoning information.</li> </ul> <p>If you deploy a Storage Resource agent and a fabric probe job has not been run, the connections from the Storage Resource agents to the switches are not displayed, and all the fibre channel (FC) ports are reported as unused</p> |                                                |                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                              |

### SMI agents

In general, it is most beneficial to use SMI agents as the primary agent type for Brocade and McDATA fabrics. See Table 17 on page 325 for the fabric information that is gathered by SMI agents. The use of SMI agents enables support for heterogeneous fabrics composed of switches from Brocade and McDATA.

The following SMI agent profiles and their respective subprofiles are supported:

Table 18. Supported SMI agent profiles and subprofiles

| Profile | Subprofile                                                                                                                                                              |
|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Fabric  | <ul style="list-style-type: none"> <li>• Zone Control</li> <li>• Enhanced Zoning and Enhanced Zoning Control</li> </ul>                                                 |
| Switch  | <ul style="list-style-type: none"> <li>• Blades</li> <li>• Access Points</li> <li>• Multiple Computer System</li> <li>• Physical Package</li> <li>• Software</li> </ul> |

The level of support depends on the subprofiles registered with the SMI agent. For example, zone control support requires that the SMI agent has the Zone Control subprofile registered with it.

SMI agents also notify the Device server of changes on the fabric, such as changes to the topology or to the zoning. This notification is done through CIM indications. The CIM indications are subscribed to for a SMI agent through the CIMOM discovery job. Some extra configuration might be required set up CIM indications for environments with systems that have multiple IP addresses.

The CIM indications send specific details of a change on the fabric to the Device server. Rather than recollect the complete topology and zoning information for every CIM indication received, the Device server analyzes the CIM indication and in most cases either sends a small set of post-indication discovery requests to the SMI agent to collect information pertinent to that event, or sets certain devices as missing without any post-indication requests made to the SMI agent. For CIM Indications that represent a zoning change, the complete zoning information is recollected from the SMI agent.

### Other agents

Although the SMI agents provide the most benefits, the use of multiple agent types also has advantages, such as redundant collection of information if one type of agent fails. Also, some types of agents provide certain features the other agents do not. However, the use of multiple agent types can increase network traffic and management load on the switches.

See Table 17 on page 325 for a list of the different types of jobs that collect data for the fabrics. You can use this table to determine which agent types to use for your fabric.

In-band, out-of-band, and Storage resource agents provide the same information, such as a change that has occurred in the SAN, by indicating that a discovery operation should be performed. The agents notify the Device server when something has happened in the SAN and then a discovery is done to identify the changes.

The fabric agents that are running on the managed hosts detect in-band and Storage Resource agent events.

The the Device server out-of-band change agent detects the SNMP traps. For the Device server to receive the SNMP traps, configure the device, such as a switch, to send traps to the Device server.

The SNMP traps that are sent by the devices and by the Device server can be sent to the IBM Tivoli Storage Productivity Center graphical user interface (GUI). The events can be displayed in the event browser, or they can be sent to another trap console for monitoring. You then can provide error recovery or other advanced processing, such as paging.

### **Fabric discovery**

This topic describes how the Fabric Manager uses SNMP queries to discover information about selected fabric switches.

Management Information Base (MIB) information is collected from those switches. Out-of-band agents are configured for the Device server and contacted from the server through SNMP.

## **Discovering computers in the Windows domain, NAS, and SAN File System**

This topic describes how a NAS discovery finds SAN File System clusters.

When you run a discovery to find Windows computers on a domain, it discovers only new Windows computers. For IBM Tivoli Storage Productivity Center to recognize a non-Windows computer, you must install an agent on that computer.

For a Windows domain, Tivoli Storage Productivity Center only performs a discovery on the first Windows system on the domain; not all the systems on the domain.

A NAS discovery finds NAS Filers, and a SAN File System discovery finds the SAN File System clusters.

For information about configuring your system for NAS and SAN File System discovery, see the section about the Manual NAS/NetWare Server Entry node in the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide*.

### **Discovering NetWare Filer systems that do not have agents**

This topic describes how you can discover NetWare Filer systems that do not have installed Storage Resource agents or Data agents.

To discover NetWare Filer systems, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services > Discovery**.
2. Right-click **NetWare Filers** and click **Run Now**.

### **Discovering Windows computers that do not have agents**

Use discovery schedules to detect Windows computers that do not have Data agents or Storage Resource agent deployed on them.

To discover Windows computers that do not have agents in your Windows domain, perform the following steps:

1. In the navigation tree, expand **Administrative Services > Discovery**.
2. Right-click **Windows Domain, NAS, and SAN FS** and click **Run Now**.
3. After you have identified the computers, you can deploy Storage Resource agent on the computers to get detailed information about the computers and file systems. Computers that are discovered by Tivoli Storage Productivity

Center but do not have Data agents or Storage Resource agents deployed on them are displayed on the Data Manager > Reporting > System-wide > Unmanaged Computers report.

## Discovering VMware VI data sources

This topic describes how you can discover VMware VI data sources.



To display information about VMware VI data sources, complete the following steps:


1. In the Navigation Tree pane, expand **Administrative Services > Discovery > VMware VI Data Source**.
2. Right-click **VMware VI Data Source**. Click **Run Now**.

## Discovering computers

This topic describes how Tivoli Storage Productivity Center gathers information from the computers that are discovered.

You must install Storage Resource agents directly on the computer it manages so that you can gather information about that computer. To gather information from the agents, you create ping, probe, and scan jobs. A scan job discovers the largest amount of information but is the most time-consuming and intensive job.

When an agent shuts down and Tivoli Storage Productivity Center detects that the agent is no longer responding, the agent is put into the missing state (a black circle  ) to indicate that it is not detectable. Tivoli Storage Productivity Center also does a heartbeat ping to the agents every 30 minutes. If this ping detects the host while the agent is not running, the host is put into the warning state (a yellow triangle  ) to indicate that the host is responding but the agent is not.

If Tivoli Storage Productivity Center can no longer communicate with an agent, the agent is marked as missing (a black circle  ). Information about the host continues to be displayed until the host is removed from the database. To remove the host from the database, you must uninstall the Storage Resource agent.

## Discover storage subsystem, tape library, and fabric information with CIM agents

This topic describes how you can discover storage subsystems, tape libraries and fabrics in your environment using CIM agents.

CIM agents can be referred to by various names, such as CIMOM (CIM Object Manager) or SMIS agent. A CIM agent consists of a CIMOM and an SMI-S provider for the managed device. CIMOM discovery performs basic health checks of the CIMOMs. Even if you do not expect the number of CIMOMs to vary, be sure to schedule a CIMOM discovery often enough to be kept informed about your environment.

To automatically discover CIMOMs that are visible to the Device server with SLP, follow the instructions in “Discover CIM agents automatically” on page 245.

To add and discover CIMOMs manually, follow the instructions in “Adding a CIM agent connection” on page 244.



Once your CIMOMs have been manually added you must perform a CIMOM discovery to gather information about the CIMOMs so that IBM Tivoli Storage Productivity Center can communicate with them.

You are not required to run discovery jobs for storage resources that connect to Tivoli Storage Productivity Center using native interfaces: IBM System Storage DS8000, SAN Volume Controller, and the XIV system.

To discover CIMOM information for any CIMOMs that you manually added, complete the following steps:

1. Expand **Administrative Services > Data Sources > CIMOM Agents**. CIMOMs will be listed in the CIMOM Agents table along with their connection status. For each automatically-discovered CIMOM you want to use, click the magnifying glass to the left of the CIMOM name in the table. The View and Edit window for that CIMOM will open. Enter missing information, such as user name and password, for that CIMOM. Also make sure you check the **Test CIMOM connectivity before updating** box. Click **File > Save** on the menu bar. This will save your information and perform a connectivity check based on the information you entered. You will need to repeat this process for each SLP-discovered CIMOM you choose to use.
2. Expand **Administrative Services > Discovery > CIMOM**. The **Edit CIMOM** page opens in the right pane.
3. Optional: Click the **Options** tab. Enter the IP addresses or host names for the directory agents to be discovered.

**Note:** If you do not want to discover every CIMOM in your system's local subnet, clear the Scan Local Subnet check box. If you clear this box the CIMOM discovery will only run on those CIMOMs already listed under **Administrative Services > Data Sources > CIMOM Agents**. If you chose not to clear this box CIMOM discovery will search for and discover all the CIMOMs in your subnet which may take a long time.

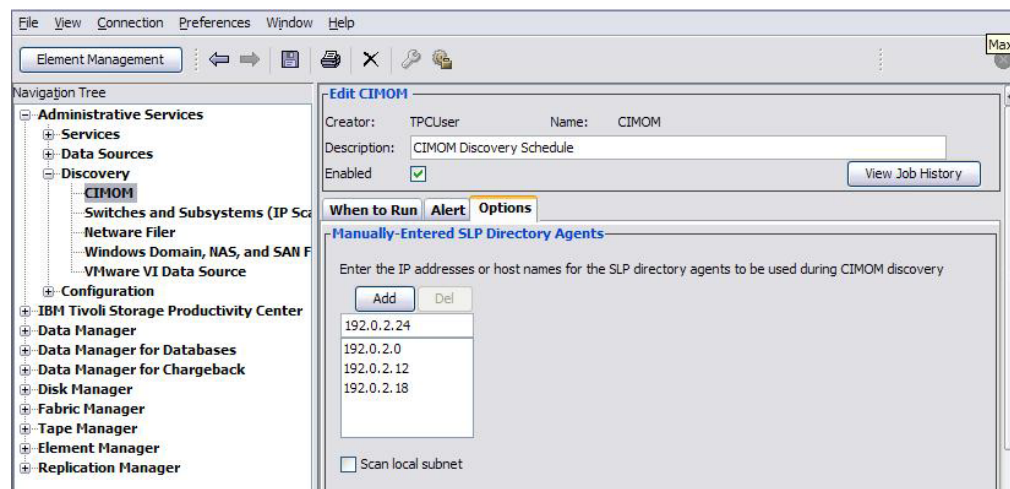


Figure 22. Editing the CIMOM discovery job (Options tab)

4. Click the **When to Run** tab and indicate how often and when to run this discovery job.
5. Click **File > Save** to save the job.

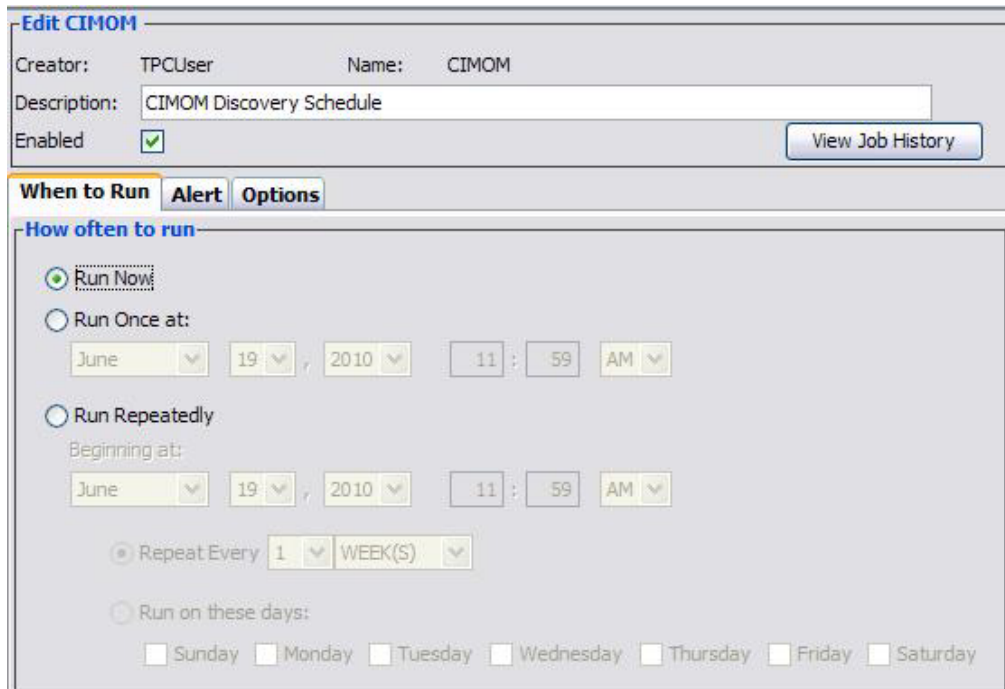


Figure 23. Editing the CIMOM discovery job (When to Run tab)

6. After you have set the parameters for a CIMOM discovery, perform the discovery to gather information on all the CIMOMs you have entered in the Options page.  
Expand **Administrative Services > Discovery**. Right-click **CIMOM** and click **Run Now**.
7. Click the Job Management icon (📁) in the toolbar to view the status of the discovery.
8. When the discovery has finished, view the list of discovered storage subsystems or tape libraries by expanding either **Disk Manager** and click **Storage Subsystems**, or **Tape Manager** and click **Tape Libraries**.



Figure 24. Checking storage subsystems

## Discovering fabrics with Storage Resource agents

This topic describes how to discover fabrics using the Storage Resource agents that are deployed on the computers in your environment.

When you deploy Storage Resource agents to computers, the agents automatically run discovery and provide information about the fabrics or switches to which the

computers are attached. You do not need to run a specific discovery step for the fabrics connected to computers on which Storage Resource agent are deployed.

For example, if you have McDATA and QLogic switches and have Storage Resource agents deployed on the computers that are attached to these switches, the switches are automatically discovered when you deploy the Storage Resource agents.

If you have a McDATA switch, ensure that the switch has the Open Systems Management Server (OSMS) feature enabled. OSMS is required for in-band discovery. The host control option must be enabled if you want to do zone control and if no CIM Agent for the fabric is present. Refer to your McDATA documentation for information about these features and how to enable them.

In-band discovery with Storage Resource agents provides the following information:

- Fabric WWN, information about fabric-to-switch relationships, and key attributes
- Host and device information (collected from Storage Resource agents only), including HBA information
- Basic information to identify the host and devices in the fabric

After discovery has run, you can run probes to collect topology and zoning information for the fabrics.

**Note:** In Tivoli Storage Productivity Center v4.2 and above, you can deploy Storage Resource agents only. If you want to install an inband fabric agent (Fabric agent) for fabric discovery and monitoring, you must own a previous version of the product. See the documentation for the previous version of the product for information about how to install and use a Fabric agent.

The vendors of fabric switches can help you determine whether to use CIM agents or Storage Resource agents to discover and monitor fabrics. For example, for Brocade and McDATA switches, use CIM agents to collect all fabric data and enable performance monitoring and zone control. Only use Storage Resource agents to collect information about HBAs or if data collection through the inband data path is preferred. For Cisco and QLogic switches, you must have one or more Storage Resource agents on the fabric to collect full topology and zoning information, and to enable the zone control functions.

See the *Planning for Fabric Manager* section in the *Tivoli Storage Productivity Center Information Center* or *Installation and Configuration Guide* for more information about using CIM and Storage Resource agents to collect information about fabrics.

## Discovering storage subsystems and switches

To find switches and storage systems in your environment, use the Switches and Subsystems (IP Scan) discovery schedule to find switches and storage systems in your environment. This discovery schedule searches for devices with a range of IP addresses that you specify.

Keep in mind the following considerations before running this discovery schedule:

- Switches and storage systems that require a CIM object manager (CIMOM) are not discovered using an IP scan. Use CIMOM discovery to discover these devices.

- Before a switch can be discovered, you must have an out-of-band fabric agent installed on the computer where the switch is installed or have the agent installed on a computer that has access to the switch.
- Before you can discover storage devices, ensure that port 7 (ECHO port) is open on the devices and Internet Control Message Protocol (ICMP) is available.
- If no CIM Agents for your Brocade switches are configured, you must discover out-of-band Fabric agents to get zone information for the Brocade switches.

To edit a discovery schedule for scanning a range of IP addresses, complete these steps:

1. In the navigation tree, expand **Administrative Services > Discovery**.
  2. Click **Switches and Subsystems (IP Scan)**. The Edit Switches and Subsystems (IP Scan) panel is displayed.
  3. Click **Options** to display the Options page.
  4. Specify the IP address ranges that you want to scan in the **From** and **To** fields. You can enter an IPv4 or IPv6 address depending on what is supported within your environment:
    - Enter IPv4 addresses in the following format: ###.###.###.###. For example, if you enter **192.168.1.68** in the **From** field and **192.168.1.254** in the **To** field, Tivoli Storage Productivity Center scans IP addresses that fall in the range between 192.168.x.68 and 192.168.x.254.
    - Enter IPv6 addresses according to the standards in your environment. The preferred IPv6 address representation is written as eight groups of four hexadecimal digits: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx, where each x is a hexadecimal digit representing 4 bits. For example: 2001:DB8:0000:1234:0000:0000:5678:ABCD. You can also specify IPv6 addresses using shortened formats that omit leading zeros 2001:DB8:0:1234:0:0:5678:ABCD or use double colons 2001:DB8:0000:1234::5678:ABCD in place of a series of zeros. For example, if you enter **2001:DB8::0** in the **From** field and **2001:DB8::FF** in the **To** field, Tivoli Storage Productivity Center scans the IP addresses that fall in the range between 2001:DB8:x:0 to 2001:DB8:x:FF. The address range that you specify determines how long an SNMP scan takes to run. For example, if the range includes over 256 addresses, the scan might take some time to complete. A warning message is displayed if you indicate a range that includes between 257 and 4096 (0xFFF) addresses. This warning message notifies you that the scan requires significant time to run as more addresses are included. Keep in mind that the scan range you enter cannot exceed 4096 (0xFFF) addresses.
- Note:** It takes about 5 minutes to search a subnet. For out of band fabric discovery to work, the FC Management MIB and FE MIB must be enabled on the switch.
5. On the **When to Run** page, determine when and how often to run the discovery.
  6. On the **Alert** page, specify any alerts to trigger if the discovery schedule fails to run.
  7. Click **File > Save** to save the discovery. The discovery runs according to the schedule that you defined on the **When to Run** page.

### Gathering switch performance data

This topic describes how to gather performance data from a switch to evaluate the switches that are in your environment.

To complete the discovery process for switch performance functions, a SMI agent must be enabled for your switch.

- For some switches, if you have the prerequisite firmware, you must enable the SMI agent (also referred to as an embedded SMI agent).
- For other switch vendors, you must install the SMI agent on any host computer and configure the agent to run providers for one or more switches. An SMI agent that is installed on a host is referred to as an external SMI agent.

Contact your switch vendor for information about how to enable the embedded SMI agent or to install the external CIM agent. For the SMI agents that are supported by IBM Tivoli Storage Productivity Center, see the support matrix at [http://www.ibm.com/support/docview.wss?rs=40&q1=switch&uid=swg21384219&loc=en\\_US&cs=utf-8&lang=en](http://www.ibm.com/support/docview.wss?rs=40&q1=switch&uid=swg21384219&loc=en_US&cs=utf-8&lang=en).

To gather switch performance data, complete the following steps:

1. Run a CIMOM discovery for a switch SMI agent.
2. Run a Fabric Probe for the Fabric that the switch is a part of.

For information about creating a switch performance monitor, see “Creating a switch performance monitor” on page 375.

---

## Collecting data

Create data collection jobs to gather information about the storage resources in your environment. There are different data collection jobs within IBM Tivoli Storage Productivity Center depending on the type of information is collected and from which types of resources the information is retrieved: discovery, probes, Tivoli Storage Productivity Center Server probes, scans, and pings.

### Collecting storage statistics (probes)

Use probes to collect statistics about the monitored storage resources in your environment, such as computers, storage subsystems, fabrics, volumes, disk controllers, hard disks, and filesystems. Probe jobs can also discover information about new or removed disks and filesystems. You can view the information collected by probes in Asset, Capacity, Storage Subsystems, and System reports.

You can define any number of probes to collect statistics for different storage resources at different times.

- From the **IBM Tivoli Storage Productivity Center** node you can collect statistics on computers, disk controllers, hard disks, clusters, fabrics, storage subsystems, volumes, tape libraries, and filesystems.
- From the **Data Manager for Databases** node you can itemize and create an inventory of the files, instance, logs, and objects that make up your enterprise's monitored RDBMSs.

**Note:** The login ID that Data Manager for Databases uses to log into Microsoft SQL Server instances that you want to probe must have "Permit" access. For information about how to set the logins for Microsoft SQL Server instances, see the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide*.



**Note:** To ensure the accuracy and consistency of the data that appears in the reports, run regularly scheduled probes against the hosts that use/import volumes

and the storage subsystems upon which you want to report. You should run probes against storage subsystems after running them against any hosts.

## Creating probes

Create and schedule a probe to job to gather information about your storage resources. You can define any number of probes to gather information about different storage resources at different times.

To create a probe, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center > Monitoring > Probes**. For Data Manager for Databases, expand **Data Manager for Databases > Monitoring > Probes**.
2. Right-click **Probes** A menu is displayed.
3. Select **Create Probe**.
4. In the **What to Probe** or **Instance** page, perform the following steps:
  - a. In the **Available** list box, highlight the resources you want to probe.
  - b. Click . The resources are moved to **Current Selections** list box. Use  to move resources to the **Available** list box.
5. In the **When to Run** page, schedule when you want the probe to run. You have the following options:
  - Run the job immediately
  - Run the job once at a specified time and date. Use the list boxes to specify the month, day, year, and the hour, minute, and AM or PM.
  - Run the job repeatedly according to a schedule you specify. You can perform the following tasks:
    - Specify the month, day, year, and the hour, minute, and AM or PM when the job is to run.
    - Indicate how often the job runs. You can use the list boxes to specify the number of minutes, hours, days, weeks, or months.
    - Indicate which days on which run the job. A check mark is displayed next to the days you select.

Use the fields in the **How to handle timezones** section to indicate the time zone that Data Manager for Databases uses when scheduling and running the action. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:

- 9:00 AM of the time zone where the server resides
- 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on a Storage Resource agent or Data agent located in the PST time zone.

### Use the timezone that the server runs in


Select this option to use the time zone of the location where the Data server resides.

### Use this timezone

Select this option to indicate a specific time zone for when to run the job. When you select this option, the list box to the right of the field activates.

6. In the **Alert** page, set the following options:
  - The condition that triggers an alert. The default setting is **Probe Failed**. No other conditions are currently available.

- The notification method of an alert: SNMP Trap, IBM Tivoli Enterprise Console Event, Login Notification, Windows Event Log, UNIX Syslog, running a script, and email. If you select SNMP Trap, Tivoli Enterprise Console Event, or email for the method by which you are notified of an alert, you must define a destination for the method on the **Administrative Services** → **Configuration** → **Alert Dispositions** window.

7. Click  to save the probe. You must save the probe before Tivoli Storage Productivity Center can collect data about storage resources.
8. Enter a name for the probe and click **OK**. An entry for the probe is displayed under the **Probes** node in the navigation tree. The probe is run at the date and time specified in the schedule definition.
9. Use the Job Management panel to manage view information about the runs and jobs that are related to the probe. For example, use the Job Management panel to view the status of a run for a schedule, the date and time when it was run, and the log files of the jobs associated with the run.

**Related reference:**


“Managing schedules, jobs, and runs” on page 295

Use the Job Management panel to view and manage all the schedules, runs, and jobs related to the storage devices that are monitored by IBM Tivoli Storage Productivity Center.

## Editing probes

Change the settings in an existing probe job.

To edit an existing probe, complete the following steps:


1. Expand **IBM Tivoli Storage Productivity Center** > **Monitoring** > **Probes**. For Data Manager for Databases, expand **Data Manager for Databases** > **Monitoring** > **Probes**.
2. Click the name of the probe you want to edit. The content pane displays information about that job.
3. Edit the probe definition.
4. Click  to save the updated probe definition.


## Creating fabric probes

Use probes to collect statistics about fabrics and fabric groups.

After you have defined your fabric groups, you can use Fabric Manager to run probes against fabrics and fabric groups to collect information.

To run probes against fabrics, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** → **Monitoring**, right-click **Probes**, and select **Create Probe**.
2. In the **What to Probe** tab enter a description for the probe, and from the **Available** column, select the fabrics and fabric groups to be probed. Move them to the **Current Selections** column by clicking .
3. In the **When to Run** tab specify when the probe is to begin. You can begin the probe immediately, schedule it to begin at a later time, and specify whether to run the probe repeatedly.
4. In the **Alert** tab specify the action to be taken if the probe fails. See “Running scripts with alerts” on page 507 for a description of these triggered actions.

5. Click  to save the probe. You must save the probe before Tivoli Storage Productivity Center can collect data about storage resources.
6. Enter a name for the probe and click **OK**. An entry for the probe is displayed under the **Probes** node in the navigation tree. The probe is run at the date and time specified in the schedule definition.
7. Use the Job Management panel to manage view information about the runs and jobs that are related to the probe. For example, use the Job Management panel to view the status of a run for a schedule, the date and time when it was run, and the log files of the jobs associated with the run.

**Related reference:**

“Managing schedules, jobs, and runs” on page 295


Use the Job Management panel to view and manage all the schedules, runs, and jobs related to the storage devices that are monitored by IBM Tivoli Storage Productivity Center.

### Creating tape library probes

Probes collect data from the tape libraries that have been discovered.

You can create tape library probes that collect information from all your tape libraries, or a subset. Probes can be scheduled to run immediately, or at a scheduled point in time. The data that probes collect is put into the database and is displayed in various places throughout the product, such as the topology viewer and the **Tape Manager > Tape Libraries** node.

To create a tape library probe:

1. Expand **IBM Tivoli Storage Productivity Center > Monitoring > Probes** and select **Create Probe**
2. Enter the information in the three tabbed pages:
  - **What to Probe** - select the tape libraries and tape library groups to include in this probe
  - **When to Probe** - specify the scheduling information for this probe
  - **Alert** - specify the triggering condition that will cause an alert to be generated for this probe, and the actions taken when an alert is generated
3. Click  to save the probe. You must save the probe before Tivoli Storage Productivity Center can collect data about storage resources.
4. Enter a name for the probe and click **OK**. An entry for the probe is displayed under the **Probes** node in the navigation tree. The probe is run at the date and time specified in the schedule definition.
5. Use the Job Management panel to manage view information about the runs and jobs that are related to the probe. For example, use the Job Management panel to view the status of a run for a schedule, the date and time when it was run, and the log files of the jobs associated with the run.

There are some special situations to note when creating and running probes of tape libraries:

- Do not include more than three tape libraries in the same probe if they are all registered with the same IBM SMI-S Agent for Tape. Doing so can increase the load on the SMI-S agent and can increase the possibility of timeouts. To avoid timeouts, create multiple probes with different start times, and split the libraries up among the probes.



- After the list of tape libraries that are registered with the IBM SMI-S Agent for Tape has been altered, a condition can occur in which not all of the tape cartridges are returned to the CIM client, in this case, Tivoli Storage Productivity Center. To resolve this situation, restart the IBM SMI-S Agent for Tape. For instructions, see the documentation for the SMI-S Agent.
- When the probe of a tape library successfully completes, some of the buttons in the Tape Libraries GUI panel might be disabled. To enable the buttons, close and then redisplay the panel.

**Related reference:**

“Managing schedules, jobs, and runs” on page 295

Use the Job Management panel to view and manage all the schedules, runs, and jobs related to the storage devices that are monitored by IBM Tivoli Storage Productivity Center.

**Locating a tape cartridge:**

View a list of all the cartridges for each discovered tape library.

Tape libraries can contain many cartridges (sometimes called tape volumes), and it can be difficult to find a particular tape cartridge among all the cartridges for a particular tape drive. IBM Tivoli Storage Productivity Center simplifies this task by listing all the cartridges for each discovered tape library. To find a particular tape cartridge:

1. Expand the **Tape Manager** node and click **Tape Libraries**.
2. Select a tape library from the list of discovered tape libraries.
3. Click **Cartridges**.
4. Click the **Label** column header to sort the list of cartridges by label. The triangle next to the label header indicates ascending or descending order.
5. Scroll through the list of cartridges to find the one you are looking for.
6. The **Location** column will indicate exactly where the tape is located in the library.

**Note:** You can also sort by other columns by clicking the column headers. For example, if you know that the cartridge you are looking for is dual sided, you can sort the list by that criterion. This can make finding a particular cartridge much easier, especially for libraries that have many cartridges.

**Creating a VMware VI probe**

Probes collect data from the VMware VI data sources that have been discovered.

To create a VMware probe, complete the following steps:

1. In the Navigation Tree pane, expand **IBM Tivoli Storage Productivity Center > Monitoring > Probes**.
2. Right-click **Probes** and click **Create Probes**.
3. Enter the information in the three tabbed pages:

**What to Probe**

Enter a description for the probe and from the **Available** column, select **Hypervisors** to include in the group to be probed. Click >> to move the group to the **Current Selections** column.

### When to Run

Specify when the probe is to begin. You can begin the probe immediately, schedule it to begin at a later time, and specify if it should be run repeatedly.

**Alert** Specify the triggering condition that will cause an alert to be generated for this probe, and the actions that should be taken when an alert is generated.

4. To save the probe, click **Save**.
5. In the **Save As** window, enter the probe name and click **OK**.

## Collecting storage statistics from Tivoli Storage Productivity Center servers (TPC server probes)

Define and run TPC server probes from the master server to collect storage information that has been gathered by subordinate servers. The information gathered by probes of subordinate servers is displayed in the **Tivoli Storage Productivity Center > Rollup Reports** node of the master server's navigation tree.

A subordinate server is a server that monitors and reports on storage resources like a standard server, but also communicates with the master server during TPC server probes. During these probes, the master server collects the storage information gathered by a subordinate server's agents and stores that information in its own database repository.

**Note:** If the same storage entity is managed by multiple subordinate servers, rollup reports reflect the storage information from the subordinate server that most recently probed that entity.

Before you run TPC server probes, make sure to perform the following:

1. Associate subordinate servers with the master server on the **Administrative Services > Data Sources > TPC Servers** window. See "Adding an Tivoli Storage Productivity Center Server" on page 280 for information on how to associate subordinate servers with a master server.
2. Run discoveries, probes, and scans on the subordinate servers to gather information about their locally-managed entities.

The information collected by TPC server probes from subordinate servers includes the following:

- Asset information: detailed statistics about agents (Tivoli Storage Productivity Center), computers, storage subsystems, disk and volume groups, disks, filesystems, logical volumes, volumes, and fabrics.
- Database asset information: detailed statistics about the RDBMSs in your environment, including Oracle, SQL Server, Sybase, and UDB.
- Capacity information: storage metrics related to the disk capacity, filesystem capacity, filesystem used space, and filesystem freespace of the storage entities in your environment.
- Database capacity information: storage metrics related to the storage capacity of the RDBMSs in your environment, including Oracle, SQL Server, Sybase, and UDB.

Additional storage information that is collected by a subordinate server and used within Tivoli Storage Productivity Center functions is available for only that subordinate server. The Tivoli Storage Productivity Center functions include the topology viewer, data path explorer, volume provisioning, volume performance, or


SAN Planner. The storage information is not collected by the master server and thus not rolled up into the reports that are available through the master server.

Also, the Tivoli Storage Productivity Center functions within the master server will use the storage information collected by the master server only.

### Creating Tivoli Storage Productivity Center server probes

Use Tivoli Storage Productivity Center server probes from the master server to collect storage information that has been gathered by subordinate servers.

To create a Tivoli Storage Productivity Center server probe, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center > Monitoring > TPC Server Probes**.
2. Right-click **TPC Server Probes** and select **Create TPC Server Probes** from the menu.
3. Enter a brief description of the probe in the **Description** field.
4. In the **What to Probe** page, perform the following steps:
  - a. In the **Available** list box, highlight the Tivoli Storage Productivity Center servers you want to probe. You can do any of the following: 1. Select the **All TPC Servers** item to automatically include all subordinate servers in the probe. 2. Select the name of a specific subordinate server and all its related storage entities that you want to probe. 3. Select the specific entities on a subordinate server that you want to probe. For example, you can expand the node for a subordinate server and select the specific entities for which you want to collect data: All Resources, Clusters, Computer Groups, Computers (including Hypervisors), Database Instances, Fabrics, RDBMS Computer Groups, Storage Subsystem Groups, and Storage Subsystems.
  - b. Click  to include highlighted Tivoli Storage Productivity Center servers or storage entities in the probe. The entities that you selected are moved to **Current Selections** list box.
5. In the **When to Run** page, schedule when you want the probe to run. You have the following options:
  - Run the job immediately
  - Run the job once at a specified time and date. Use the list boxes to specify the month, day year, and the hour, minute, and AM/PM.
  - Run the job repeatedly according to a schedule you specify. You can perform the following tasks:
    - Specify the month, day, year, and the hour, minute, and AM or PM when the job is to run.
    - Indicate how often the job runs. Use the list boxes to specify the number of minutes, hours, days, weeks, or months.
    - Indicate which days on which run the job. A check mark is displayed next to the days you select.

Use the fields in the **How to handle timezones** section to indicate the time zone that Data Manager for Databases uses when scheduling and running the action. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:

- 9:00 AM of the time zone where the server resides


- 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on a Storage Resource agent or Data agent located in the PST time zone.

**Use the timezone that the server runs in**

Select this option to use the time zone of the location where the Data server resides.

**Use this timezone**

Select this option to indicate a specific time zone for when to run the job. When you select this option, the list box to the right of the field activates.

6. In the **Alert** page, set the following options:
  - The condition that triggers an alert. The default setting is **Probe Failed**. No other conditions are currently available.
  - The notification method of an alert: SNMP Trap, IBM Tivoli Enterprise Console Event, Login Notification, Windows Event Log, UNIX Syslog, running a script, and email. If you select SNMP Trap, Tivoli Enterprise Console Event, or email for the method by which you are notified of an alert, you must define a destination for the method on the **Administrative Services** → **Configuration** → **Alert Dispositions** window.
7. Click  to save the probe. You must save the probe before Tivoli Storage Productivity Center can collect data about storage resources.
8. Enter a name for the probe and click **OK**. An entry for the probe is displayed under the **TPC Server Probes** node in the navigation tree. The probe is run at the date and time specified in the schedule definition.
9. Use the Job Management panel to manage view information about the runs and jobs that are related to the probe. For example, use the Job Management panel to view the status of a run for a schedule, the date and time when it was run, and the log files of the jobs associated with the run.

**Related reference:**


“Managing schedules, jobs, and runs” on page 295

Use the Job Management panel to view and manage all the schedules, runs, and jobs related to the storage devices that are monitored by IBM Tivoli Storage Productivity Center.

**Editing TPC Server probes**

Learn how to edit the definition of a TPC server probe.

To edit an existing TPC server probe, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** > **Monitoring** > **TPC Server Probes**.
2. Click the name of the TPC server probe you want to edit. The content pane displays information about that probe.
3. Edit the values on the What to PROBE, When to Run, and Alert pages to change the TPC server probe definition.
4. Click  to save the updated TPC server probe definition.

**Managing the runs and jobs for TPC Server Probes**

View information about the runs, jobs, and log files for TPC Server Probes.

Tivoli Storage Productivity Center creates a run each time a TPC Server Probe schedule collects data from subordinate servers. A job is a unit of work within a run for a schedule. Use the Job Management panel to manage TPC Server Probe schedules and view information about the related runs and jobs. Specifically, use this panel to perform the following tasks for schedules:


- View or edit the definition of a schedule.
- Run a schedule immediately.
- Rename or delete a schedule.
- View basic information about a schedule, including the ID of the user who created it, description, the last time the schedule was run, and the next time the schedule is run.

Use the Job Management panel to perform the following tasks for the runs and jobs related to TPC Server Probe schedules:

- View the status of a run and the date and time when it was executed.
- View information about the jobs associated with each run, including the name of the computer on which a job was run, the type of the monitored device, and the status of the job.
- View the log files for one or more jobs at a time.
- Display jobs that have started or completed within a specified length of time. For example, display jobs that have completed in the last two weeks.
- Display jobs that have a certain status, such as Success, Running, or Failed.

To access the Job Management panel for a TPC Server Probe schedule, complete the following steps:

1. In the navigation tree, expand **IBM Tivoli Storage Productivity Center > Monitoring > TPC Server Probes**.
2. Select the probe schedule whose runs and jobs you want to manage. The content pane displays information about the schedule.
3. In the content pane, click **View Job History**. The Job Management panel is displayed and the selected schedule is highlighted in the **Schedules** section. View the entries in the **Jobs for Selected Schedules** section to manage the jobs and runs for the schedule.

You can also access the Job Management panel at any time by clicking the related icon () in the toolbar. Use the **Filters** section on the Job Management panel to display the TPC Server Probe schedule that you want to manage. See “Managing schedules, jobs, and runs” on page 295 for more information.

## Collecting storage usage and trending information (scans)

Use scans to collect statistics about the usage and trending of storage consumption of storage resources in your environment. You can view the information collected by scans in reports and use it as part of a quota analysis.

Use a Data agent or Storage Resource agent to collect information for a scan. These agents must be deployed on the storage resources that you want to monitor. By running scans and generating the results of those scans in reports, you can perform the following tasks:

- View when files were created, accessed, and modified and by what group or user.
- Map storage resources to the consumers of those resources.

- Generate historical views of storage consumption and utilization that show usage trends over time. System administrators can use these views to project and plan for future storage consumption.
- Collect detailed RDBMS space usage reports at the table, index, and user levels.
- Gather individual user space usage at the table space, database, instance, computer, and network wide levels.
- Collect detailed space usage information about monitored tables and their associated indexes. For example, when scanning UDB tables, this information includes many of the calculated REORGCHK values. These include Used Page %, Used Space %, and Cluster Ratio.
- Historically track database, table space, and user usage and use that data to forecast and display future usage.
- Gather information for quotas that can be created to trigger alerts (and create reports) when users exceed a specified amount of space usage.
- (UDB only) Perform (optionally) runstats on all tables involved in the scan process.
- Assign costs to space used at the user level when combined with Data Manager for Chargeback.

**Note:**


1. A licensed agent must be installed on a computer that contains instances against which you want to run a scan. You must run a probe against an instance before you can select any tablespaces against which you want to run a scan.
2. The login ID that Data Manager uses to log in to Microsoft SQL Server instances that you want to scan must have "Permit" access.

### Creating scans

Create and schedule a scan to gather usage information about your storage resources. You can specify which file systems, computers, clusters, directories, databases, and other storage resources that you want to scan. You can define any number of scans to gather information about different storage resources at different times.

To create a scan, complete the following steps:

1. Expand **Data Manager > Monitoring > Scans**. For Data Manager for Databases, expand **Data Manager for Databases > Monitoring > Scans**.
2. Right-click the **Scans** node and select **Create Scan**.
3. For Data Manager, define the following scan properties:
  - The computers and file systems whose storage usage and consumption you want to monitor. For more information, see “Including storage resources in a scan” on page 345.
  - The directory groups to scan. For more information, see “Selecting directory groups to scan” on page 346.
  - The profiles to use during a scan. For more information, see “Using profiles in a scan” on page 348.
  - The schedule when you want a scan to run. For more information, see “When to run the scan” on page 348.
  - An alert to notify you if the scan fails. For more information, see “Defining alert values for a scan” on page 349.
4. For Data Manager for Databases, define the following scan properties:

- The databases, table spaces, instances, and computers whose storage usage and consumption you want to scan. For more information, see “Selecting databases, tablespaces to scan” on page 347.
  - The table groups you want to scan. For more information, see “Select the table groups to scan” on page 347.
  - The profiles to use during a scan. For more information, see “Using profiles in a scan” on page 348.
  - The schedule when you want a scan to run. For more information, see “When to run the scan” on page 348.
  - The options for Oracle and UDB. For more information, see “Specify options for Oracle and UDB” on page 349.
  - An alert to notify you if the scan fails. For more information, see “Defining alert values for a scan” on page 349.
5. Click  to save the scan. You must save the scan before Tivoli Storage Productivity Center can collect data about storage resources.
  6. Enter a name for the scan and click **OK**. An entry for the scan is displayed under the **Scans** node in the navigation tree. The scan is run at the date and time specified in the schedule definition.
  7. Use the Job Management panel to manage view information about the runs and jobs that are related to the scan. For example, use the Job Management panel to view the status of a run for a schedule, the date and time when it was run, and the log files of the jobs associated with the run.

**Related reference:**

“Managing schedules, jobs, and runs” on page 295

Use the Job Management panel to view and manage all the schedules, runs, and jobs related to the storage devices that are monitored by IBM Tivoli Storage Productivity Center.

**Including storage resources in a scan:**

In a scan, include the computers, computer groups, file systems, file system groups, and clusters whose storage usage and consumption that you want to monitor.

Use the **Filesystems** page to select the storage resources that you want to scan. Include clusters, individual file systems and computers, or file system and computer groups (that have been predefined using the “Working with groups of resources” on page 371 facility).


The **Current Selections** list box shows the storage resources that have been selected for the scan. The **Available** list box shows which storage resources are still available for selection in the scan.

**Note:** A cluster name is displayed next to individual computers that are members of a PowerHA or Microsoft Cluster Server cluster. Cluster resource groups are identified with a special icon. Clusters displayed in the **Available** list box cannot be expanded to show member nodes. You can move a cluster to the **Current Selections** list box to view and, if necessary, exclude member nodes.


**To include storage resources in a scan:**

1. Highlight the file systems, file system groups, computer groups, computers, and clusters you want to include in the scan in the **Available** list box.

**Note:** Before you can include file systems in a scan, you must probe the computers that contain those file systems. If you do not probe a computer before you define a scan, the file systems on that computer are unavailable for the scan.

2. Click . The storage resources that you selected are displayed in the **Current Selections** list box.

**To remove storage resources from a scan:**

- Highlight the file systems, file system groups, computer groups, computers, and clusters that you want to remove and click . This action removes the storage resources from the **Current Selections** list box and moves them to the **Available** list box.

**To exclude one or more file systems or computers from a selected cluster or group:**

Perform the following steps if you want to exclude specific file systems or computers from a cluster or group that you included in a scan:

1. Double-click the cluster or group name in the **Current Selections** list box.
2. Right click the file systems or computers that you want to exclude. A pop-up menu is displayed.
3. Select **Exclude**. This action creates an **Excludes** section in the **Current Selections** list box that displays the names of the file systems or computers that you have excluded from a scan. This action is useful when you want to select a group or cluster for scanning, but you want to exclude a few file systems or computers from that monitoring.


**Selecting directory groups to scan:**

This topic describes how to select the directory groups that you want to scan.


Use the Directory Groups tab to select the groups of directories that you want to scan. Including directory groups within a scan enables you to monitor and gather information about specific directories across a number of machines.

The **Selected Directory Groups** list box shows what directory groups have been included in the scan. The **Available Directory Groups** list box shows the groups that are still available for selection. You can select directory groups that have been predefined using the groups facility.

**To select a directory group:**

1. Highlight the desired directory group in the **Available Directory Groups** list box.
2. Click . The group you selected appears in the **Selected Directory Groups** list box.

**To remove a directory group from a scan:**

- Highlight the groups you want to remove and click . This will remove the groups from the **Selected Directory Groups** list box and move them to the **Available Directory Groups** list box.



### To create a new directory group:

- Click **New Directory Group**. The **New Directory Group** window opens, enabling you to create a new group. See “Creating directory groups” on page 372 for more information on how to create a directory group.


### Selecting databases, tablespaces to scan:

This topic describes how to select the databases and tablespaces that you want to scan.

Use the Database-Tablespaces page to select what databases, tablespaces, instances, and computers whose storage usage and consumption you want to monitor. You can include single or groups of databases, tablespaces, and computers against which to run the scan.


The **Current Selections** list box shows the databases, tablespaces, tablespace groups, computers, and computer groups that have been selected for the job. The **Available** list box shows what items are still available for selection.

### To include a database, tablespace, computer, instance, or group in a scan:

1. Highlight the desired databases, tablespaces, computers, instances, or groups in the **Available** list box.
2. Click . The databases, tablespaces, computers, instances, or groups you selected appear in the **Current Selections** list box.

For example, if you select a computer to include in a scan, that job will scan all the databases and/or tablespaces on all the instances that reside on that computer.

### To remove a database, tablespace, computer, instance, or group from a scan:

- Highlight the items you want to remove and click . This will remove the databases, tablespaces, computers, instances, or groups from the **Current Selections** list box and move them to the **Available** list box.

### Select the table groups to scan:


Select the table groups that you want to scan.

Use this page to select the table groups you want to scan. A table group represents a group of defined tables. By choosing a table group on this window, you can perform more detailed scans of tables that might not appear within other Data Manager for Databases reports (such as the Largest Table or Monitored Tables reports).


**Note:** Before you can view detailed information about tables in an RDBMS, you must have 1) defined a table group that contains the tables about which you want to view storage information and 2) scheduled and run a scan against that table group.

The **Selected Table Groups** list box shows the table groups that have been selected for the job. The **Available Table Groups** list box shows the table groups that are still available for selection.

### To include a table group in a scan:

1. Highlight the desired group in the **Available Table Groups** list box.
2. Click . The group you selected appears in the **Selected Table Groups** list box.

**To remove a table group from a scan:**

- Highlight the groups you want to remove and click . This will remove the groups from the **Selected Table Groups** list box and move them to the **Available Table Groups** list box.

**To create a new table group:**



- Click **New Table Group**. The **New Table Group** window opens, enabling you to create a new group.

**Using profiles in a scan:**

Select the profile to use as part of a scan. Profiles enable you to control what files are scanned and what statistics are gathered.


Use the **Profiles** page to select which profiles to use during a scan. Using profiles is a powerful method for determining the statistics you want to gather during scans. Keep in mind that the more profiles you include within a scan, the longer the scan will take to gather the statistics specified by those profiles.

**To include profiles in a scan:**

1. Highlight the desired profiles in the **Available Profiles** list box.
2. Click the top  to apply the selected profiles to file systems; click the bottom  to apply the selected profile(s) to directories. The profiles you selected appear in the either the **Profiles to apply to Filesystems** or **Profiles to apply to Directories** list box

**Note:** To collect statistics according to Data Manager default settings, select all of the default profiles.

**To remove profiles from a scan:**

- Highlight the profiles you want to remove and click the corresponding . This will remove the profile from the list boxes on the right and move them to the **Available Profiles** list box.

**To create a new profile:**

- Click **New Profile**. The **Statistics** tab on the **Create Profile** window is displayed, enabling you to create a new profile. See “Using profiles” on page 352 for more information.

**Related concepts:**

“Using profiles” on page 352

Use profiles to specify what statistical information is gathered and to fine tune and control what files are scanned during a scan. You can also use profiles with some reports to determine what statistics are displayed.

**When to run the scan:**

Schedule when and how often to run a scan.

You have the following options for determining when to run a scan:

- Run the job immediately
- Run the job once at a specified time and date. Use the list boxes to specify the month, day, year, and the hour, minute, and AM or PM.
- Run the job repeatedly according to a schedule you specify. You can:
  - use the list boxes to specify the month, day, year, and the hour, minute, and AM or PM when the job should begin running.
  - indicate how often the job should run. You can use the list boxes to specify the number of minute(s), hour(s), day(s), week(s), or month(s).
  - select which days you wish to run the job.

Use the fields in the **How to handle timezones** section to indicate the time zone that Data Manager should use when scheduling and running the scan. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:

- 9:00 AM of the time zone where the server resides
- 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

**Use the timezone that the server runs in.** Select this option to use the time zone of the location where the Data server resides.

**Use this timezone.** Select this option to indicate a specific time zone for when to run the job. When you select this option, the list box to the right of the field will activate.

### **Specify options for Oracle and UDB:**

Specify options for Oracle and UDB when defining a scan.

Use the Options page to do the following:

- Determine if the job should scan the freelist chains associated with the segments in your Oracle environment. Free list chains represent the pages allocated to a segment that have enough free space for additional rows. Note that checking this option will increase the time required by the scan to complete.
- Set the number of pages Data Manager for Databases should scan on a chain before stopping.
- Determine if the job should gather segment extent data for locally-managed tablespaces.
- Run the DB2 utility runstats against the tables being scanned. Checking this box will enable Data Manager for Databases to collect the most up-to-date storage statistics about tables from the DB2 catalog. Note that checking this option will increase the time required by the scan to complete, but will ensure that table and user statistics are up-to-date in UDB/DB2 databases. Alerts based on table and user stats will not function until runstats is run.

### **Defining alert values for a scan:**

This topic describes how to set an alert that will be triggered if the scan job fails.

Use the Alert page to can define an alert to notify you if a scan fails and how to notify you. You can choose to be notified by SNMP trap, IBM Tivoli Enterprise

Console event, login notification, Windows event log, UNIX Syslog, running a script, and email. If you select SNMP trap, Tivoli Enterprise Console event, or email, you must also define a destination (for example, an email address). Define this destination by expanding **Administrative Services > Configuration > Alert Dispositions**.

To define an alert for a scan, complete the following steps:

1. Click the **Alert** tab; the Alert page opens.
2. In the **Triggering Condition** area, select Scan Failed. This indicates that an alert will be triggered if the scan fails to run.
3. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition. You can choose from the following actions:

**SNMP trap**

An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

**TEC / OMNIBus Event**

A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, and so on.

**Login notification**

The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

**Login ID**

The ID of the user who receives alert notification upon logging in to the system.

**Windows Event Log, UNIX Syslog**

The ID of the user who receives alert notification upon logging in to the system.

**Event Type (Windows Event Log only)**

The type of event that is recorded to the OS log.

**Facility (UNIX Syslog only)**

The location where the UNIX Syslog is sent. You can select User or Local.

**Level (UNIX Syslog only)**

The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

**Run script**

A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window are passed to the script when the alert is triggered

**E-mail** An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit e-mail** to change the actual email message.


4. Click **File > Save** to save the scan.

5. Enter a name for the scan and click **OK**. An entry for the scan appears under the **Scans** node.

## Editing scans

Change the settings in an existing scan job.

To edit an existing scan, complete the following steps:

1. Expand **Data Manager > Monitoring > Scans**. For Data Manager for Databases, expand **Data Manager for Databases > Monitoring > Scans**.
2. Click the name of the scan you want to edit. The content pane displays information about that job.
3. Edit the scan definition by changing the values that appear on the following pages:
  - Filesystems. Use this page to edit the computers and file systems whose storage usage and consumption you want to monitor as part of a scan.
  - Directory Groups. Use this page to edit the groups of directories that you want to scan.
  - Databases-Tablespaces. Use this page to edit the databases, tablespaces, instances, and computers whose storage usage and consumption you want to monitor as part of a scan.
  - Table Groups. Use this page to edit the table groups that you want to scan. A table group represents a group of defined tables.
  - Profiles. Use this page to edit which profiles are used during a scan. Using profiles is a powerful method for determining the statistics you want to gather during scans.
  - When to Run. Use this page to edit when a scan is run.
  - Options (Oracle and UDB only). Use this page to edit the scan options for Oracle or UDB.
  - Alert. Use this page to edit how you are notified if a scan fails.
4. Click  to save the updated scan definition.

## Scanning NetWare volumes

Use scans to gather detailed information about the NetWare volumes within your environment.

However, before you can gather information about NetWare volumes, you must:

1. Configure Data Manager to discover all the NetWare servers in your NDS trees and then run a discovery job. See the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* for information about configuring and running discovery against NDS trees.

**Note:** Only the Data agents support NDS trees.

2. Access the **Administrative Services → Configuration → Scan/Probe Agent Administration** window to assign agents to each of the volumes discovered on your licensed NetWare servers. You can assign any agent to a volume that has the ability to access that volume.

| Field Name | Description                              |
|------------|------------------------------------------|
| Server     | The servers discovered in the NDS trees. |
| Domain     | The domain where a server resides.       |

| Field Name       | Description                                                         |
|------------------|---------------------------------------------------------------------|
| Tree Name        | The names of NDS trees discovered by agents.                        |
| Filesystem       | The file systems/volumes discovered in the NetWare servers.         |
| Scan/Probe Agent | Agents assigned to the file systems/volumes in the NetWare servers. |

**Note:** If you do not assign an agent to a volume, Data Manager will not scan and gather information about that volume.

- Define a scan and include the NetWare volumes and filesystems for which you want to gather detailed storage information.

### Using profiles

Use profiles to specify what statistical information is gathered and to fine tune and control what files are scanned during a scan. You can also use profiles with some reports to determine what statistics are displayed.

Until you begin creating your own profiles, you can use the default profiles that are supplied with IBM Tivoli Storage Productivity Center. Each of these default profiles allows you to select a specific statistic that you want to gather. If you want to scan a computer or filesystem, and want to gather all of the statistical information available, simply select all of the default profiles when you create your scan job.

See “Default profiles” for more information about default profiles.

#### Default profiles:

Default Profiles are provided for your use with IBM Tivoli Storage Productivity Center. Each default profile represents a specific statistic that you can gather (which is included in the name of the profile for convenient identification).

If you want to scan a computer or filesystem and gather all of the statistical information available, select all of the default profiles when you create your scan. TPCUser is shown as the creator of the default profiles, and TPCUser is the prefix name for each of the profiles. The following table identifies the default profiles and describes their functions.

| Default profile name | Description                                                                                                                                                                                                                             |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                      | <b>Data Manager</b>                                                                                                                                                                                                                     |
| By Access            | Gathers statistics by length of time since last access of files.                                                                                                                                                                        |
| By Creation          | Gather storage statistics by the length of time since the creation of files. This profile applies to scans run against Windows computers only. Scans with this profile do not gather any creation time data from non-Windows computers. |
| By Mod Not Backed Up | Gathers statistics by length of time since last modification (only for files not backed up since modification). (Windows only)                                                                                                          |
| By Modification      | Gathers statistics by length of time since last modification of files.                                                                                                                                                                  |

| <b>Default profile name</b>            | <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Size Distribution</b>               | Gathers information on the size distribution of files.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Largest Directories</b>             | Gathers statistics on the <i>n</i> largest directories. (20 is the default amount.)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Largest Files</b>                   | Gathers statistics on the <i>n</i> largest files. (20 is the default amount.)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Largest Orphans</b>                 | Gathers statistics on the <i>n</i> largest orphan files. (20 is the default amount.)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Most At Risk</b>                    | Gathers statistics on the <i>n</i> files that have been modified the longest time ago and have not yet been backed up since they were modified. (20 is the default amount.) (Windows only)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Oldest Orphans</b>                  | Gathers statistics on the <i>n</i> oldest orphan files. (20 is the default amount.)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Most Obsolete Files</b>             | Gathers statistics on the <i>n</i> "most obsolete" files (i.e., files that have not been accessed or modified for the longest period of time). (20 is the default amount.)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Summary By File Type</b>            | Summarizes space usage by file types (for example, .exe., .dll, .doc, .mp3, etc.).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Summary By Filesystem/Directory</b> | Summarizes space usage by file system or directory.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Summary By Group</b>                | Summarizes space usage by OS Group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Summary By Owner</b>                | Summarizes space usage by Owner.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Temporary Files</b>                 | Gathers statistics on the non-OS files not accessed in the last year and orphan files.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Wasted Space</b>                    | Gathers statistics on the non-OS files not accessed in the last year and orphaned files.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Data Manager for Databases</b>      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>DB User Space</b>                   | Gathers statistics about the user space usage within a database or tablespace.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Largest Tables</b>                  | Gathers statistics on the <i>n</i> largest tables. (20 is the default amount)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Segment Most Extents</b>            | Gathers statistics on the <i>n</i> segments with the most extents. (20 is the default amount)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Segment Most Unused Space</b>       | Gathers statistics on the <i>n</i> segments with the most empty used space. (20 is the default amount). Empty used space represents the space allocated to a segment that is currently empty/not being used. Gathering this information can help you discover space that can be reclaimed and allocated to other objects. Specifically, use this Profile to: <ul style="list-style-type: none"> <li>• find objects that are over-allocated (not using all their associated free space)</li> <li>• find the free space that exists above a table's "high-water" mark (the "high water mark" is the highest block that was written for a table)</li> <li>• find segments that have the largest amount of blocks that are not filled</li> </ul> |
| <b>Summary</b>                         | Summarizes space usage by database or tablespace, such as how many segments are in a tablespace (Oracle only), how many of those segments are tables (Oracle only), indexes, etc.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## Creating a profile:

Create a profile for use within a scan.

To create a profile perform the following steps:

1. Expand **Data Manager > Monitoring**. For Data Manager for Databases, expand **Data Manager for Databases > Monitoring**.
2. Perform one of the following actions:
  - Select **Create Profile** from the pop-up menu.
  - Click **New Profile** from the **Profiles** tab of the Create scan window.

Use the Create Profile window to specify exactly what level of detail to collect when running a scan. The more information that you specify, the more reports that you will be able to run against the information that you collect during the scan. Statistics are tied to a specific profile.

3. The **Statistics** tab on the **Create Profile** window is displayed.
4. Select statistics you want the scan to collect. Statistics are tied to a specific profile.  
See “Default profiles” on page 352 for more information about statistics.
5. If you want to further limit specific files to be scanned, select the **File Filters** tab. You can define exclusive conditions that filter what files are analyzed during the scan.  
See the “Creating file filters” on page 357 section for more information on using this tab.
6. The first time you view a new profile, this window is blank except for the text **All files selected**. That is the default and indicates that information will be collected for all files. To create your first group or condition, simply right-click **All files selected** and choose to create a new group or condition.
7. In the **File Filters** page, the default condition **All files selected** is displayed. Right-click **All files selected** and select **New Condition** from the pop-up menu. You can also specify groups for these conditions, by selecting **New Group**.



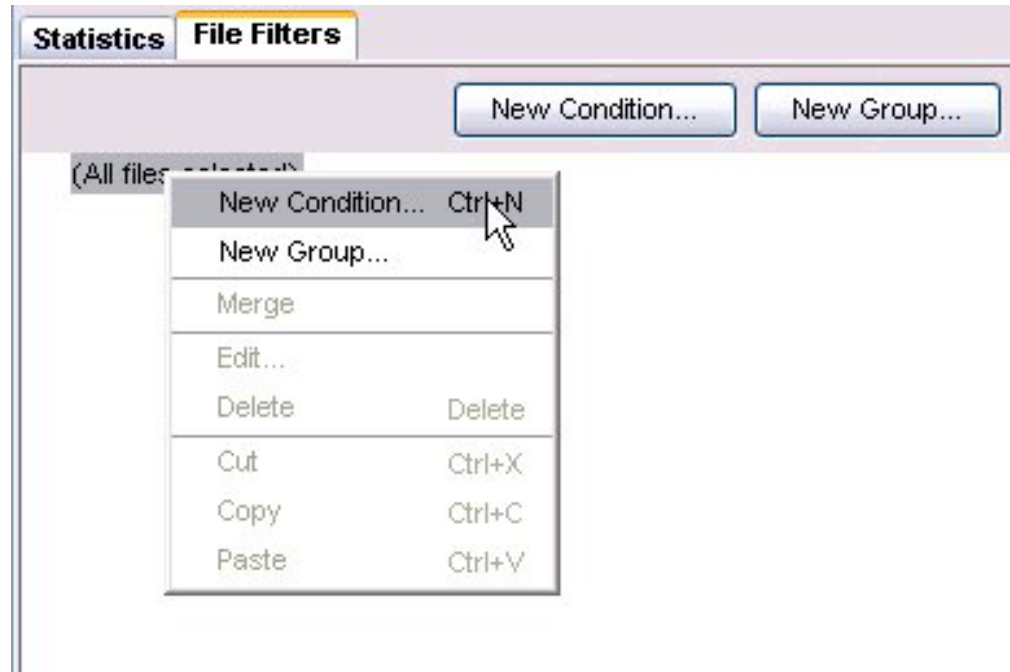


Figure 25. File Filters tab, pop-up menu

8. In the figure above, **New Condition** is selected. The **Create Condition** window opens, helping you to define your file conditions. You can create conditions based on access time, creation time, file name or group, modification time, file attributes, size, type, owner, or length.

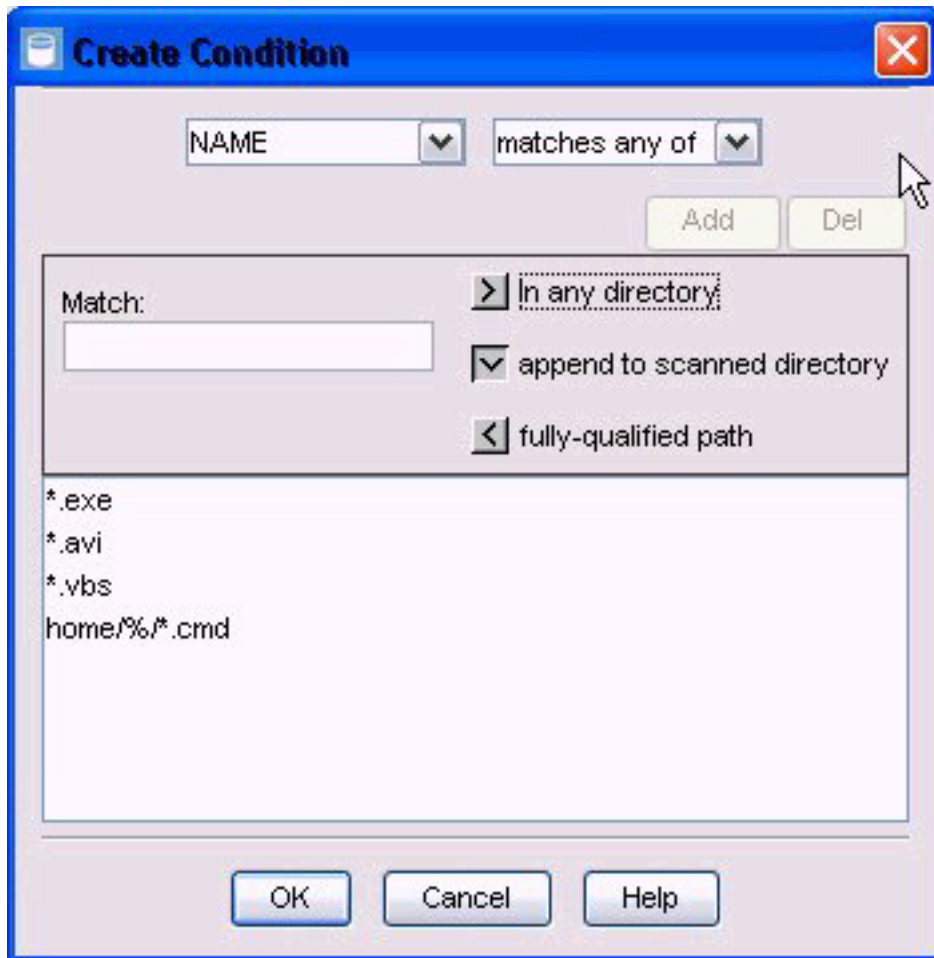


Figure 26. Create Condition window

9. Press **OK** after defining each condition.
10. If you are defining more than one condition, the **Create Group** dialog box displays to help you associate the conditions you created:





Figure 27. Create Group Window

11. Select how you want to associate the conditions you have defined. See “Associating conditions with groups” on page 362 for more information on creating groups.
12. The condition is displayed in the **File Filter** window.



Figure 28. File Filter window

13. Click  to save the profile.
14. Enter a name for the profile and click **OK**. An entry for the profile appears as a node under the **Profiles** node.

*Creating file filters:*

Use file filters as part of a profile to further define the files that you want to include or exclude in a scan. Filters limit the files on which you are gathering data during the scan based on criteria that you define.

In the following example, a group (All of) has been created with three conditions:

- The first controls the type of files to collect information on (MP3, GIF, AVI, and JPG files).
- The second specifies that the files must have been created within the last year.
- The third limits the profile to analyzing only files greater than 500 KB.

To create file filters within profile, complete the following steps:

1. To edit the file filters in an existing profile, expand **Data Manager > Monitoring > Profiles** and select an existing profile in the navigation tree. To create a new profile and add file filters, expand **Data Manager > Monitoring**, right-click **Profiles**, and select **Create Profile**.
2. Click the **File Filters** tab.
3. The first time you view a new profile, this window is blank except for **All files selected**. Selecting all files is the default and indicates that information is collected for all files. To create your first group or condition, right-click **All files selected** and choose to create a new group or condition.
4. In the File Filters page, the default condition **All files selected** is displayed. Right-click **All files selected** and select **New Condition**. You can also specify groups for these conditions by selecting **New Group**.

The **Create Condition** window opens, enabling you to define your file conditions. You can create conditions based on access time, creation time, file name or group, modification time, file attributes, size, type, owner, or length.

5. Press **OK** after defining each condition. See “Defining a filter condition” for more information on creating groups.

Each condition is displayed in the **File Filter** window.

If you are defining more than one condition, the **Create Group** dialog box helps you associate the conditions you create.

6. Associate the conditions you have defined.

See “Associating conditions with groups” on page 362 for more information on creating groups.

*Defining a filter condition:*

Define a condition for a filter.

On the File Filters page, right-click **All files selected** and select **New Condition**. The Create Condition window opens.

The first drop down list displays file criteria that you can select to define your condition. You can define file criteria based on access time, creation time, file name or group, modification time, file attributes, size, type, owner, or length. Depending on the criteria you choose to define, the second drop down list and window fields change dynamically to help you set parameters for the selected criteria. The following table describes the file criteria you can select for your file filter.

| Condition            | Description                                                                                                                                                                                                                   | Method for filtering the files                                                                                                   | Condition Modifiers                                                                                                                                                                                      |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Last Accessed</b> | Use this condition to limit files based on the last time they were accessed. For example, you can define a filter that includes files that were last accessed between 12:01 on January 1, 2005 and 11:59 on January 30, 2005. | <ul style="list-style-type: none"> <li>• later than</li> <li>• earlier than</li> <li>• between</li> <li>• not between</li> </ul> | <ul style="list-style-type: none"> <li>• Length of time before scan: enter the years, weeks, days, hours, or minutes before a scan</li> <li>• Timestamp - select files based on the timestamp</li> </ul> |
| <b>Created</b>       | Use this condition to limit files based on when they were created. For example, you can define a filter that includes files that were created between 12:01 on January 1, 2005 and 11:59 on January 30, 2005.                 | <ul style="list-style-type: none"> <li>• later than</li> <li>• earlier than</li> <li>• between</li> <li>• not between</li> </ul> | <ul style="list-style-type: none"> <li>• Length of time before scan: enter the years, weeks, days, hours, or minutes before a scan</li> <li>• Timestamp: select files based on the timestamp</li> </ul>  |

| Condition            | Description                                                                                                                                                                                                       | Method for filtering the files                                                                                                                                                                                                                         | Condition Modifiers                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Name</b>          | Use this condition to limit files based on their name. For example, you can define a filters that include files that end in *.html, or include files that end in .cmd and appear within a specified subdirectory. | <ul style="list-style-type: none"> <li>• matches any of</li> <li>• matches none of</li> <li>• matches</li> <li>• does not match</li> </ul>                                                                                                             | <p>Enter the file name upon which you want the condition to be based. For example, you can specify, *.exe, *.mp3, *.avi. in the <b>Match</b> field. See “Using wildcards in file filters” on page 363 for information about how to use wildcards (*, %, etc.) in this field.</p> <p>Select the location where the file name will be searched during the scan:</p> <ul style="list-style-type: none"> <li>• <b>In any directory.</b> During a scan, Data Manager will look for the file name in any directory. In this condition, the search string will occur at the end of the full directory path.</li> <li>• <b>Append to scanned directory.</b> This option enables you to restrict what directory to scan. Data Manager will look for the file type in the specified file system root directory or in a monitored directory in a directory group tied to a scan. In this option, you must first specify what directory groups you want scanned in the <b>Directory Groups</b> tab of the <b>Scan</b> window. See “Selecting directory groups to scan” on page 346 for more information.</li> <li>• <b>Fully-qualified path.</b> Enter the fully-qualified path when you enter the file type in the <b>Match</b> field. During a scan, Data Manager will look for the file type in the exact path you specify here. This is useful when you want to exclude system directories and files in a Most Obsolete Files report.</li> </ul> |
| <b>Group</b>         | Use this condition to limit files based on the OS group to which those files belong.                                                                                                                              | <ul style="list-style-type: none"> <li>• matches any of</li> <li>• matches none of</li> <li>• between</li> <li>• not between</li> <li>• matches</li> <li>• does not match</li> <li>• &gt;=</li> <li>• &lt;</li> <li>• &gt;</li> <li>• &lt;=</li> </ul> | Name of the OS group                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Last Modified</b> | Use this condition to limit files based on when they were last modified.                                                                                                                                          | <ul style="list-style-type: none"> <li>• later than</li> <li>• earlier than</li> <li>• between</li> <li>• not between</li> </ul>                                                                                                                       | <p>Specify the time as:</p> <ul style="list-style-type: none"> <li>• Length of time before scan: enter the years, weeks, days, hours, or minutes before a scan</li> <li>• Timestamp: select files based on the timestamp</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

| Condition         | Description                                                                                           | Method for filtering the files                                                                                                                          | Condition Modifiers                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------|-------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Attributes</b> | Use this condition to limit files based on file attributes such as read only, system, temporary, etc. | <ul style="list-style-type: none"> <li>• include any of</li> <li>• include none of</li> <li>• include all of</li> <li>• don't include all of</li> </ul> | <p>Select the file attribute upon which you want the condition to be based. You can select:</p> <ul style="list-style-type: none"> <li>• Read only: Users cannot write to the file.</li> <li>• Hidden: A file that cannot be seen by users during a normal directory listing.</li> <li>• System: A file used by the operating system that cannot normally be viewed by users.</li> <li>• Archive: A file that has been backed-up and the archive bit is turned off. An Archive bit signifies that the file has been changed since it was last archived.</li> <li>• Normal: A file with no special attributes set to it.</li> <li>• Temporary: (Windows only) A temporary or backup copy of a file.</li> <li>• Sparse: A file with a special kind of compression that does not take up a lot of CPU time. The logical size of the file looks bigger than is actually allocated. It contains some blank pages which are not on disk.</li> <li>• Reparse Point: New NTFS file system objects that have a definable attribute containing user-controlled data and are used to extend functionality in the input/output (I/O) Subsystem.</li> <li>• Compressed: A file to which a compression algorithm was applied. This file takes up less space on disk, but it will be slower to read. For Novell NetWare, the file itself is not usually compressed, rather the whole volume is compressed.</li> <li>• Offline: A file that was migrated to tape.</li> <li>• Not Content Indexed: A file that is not indexed--the Windows indexing service was told that this is not a file that should be searched through for text strings.</li> <li>• Encrypted: A file that was saved with a password (is password-protected).</li> <li>• Orphaned: (for Windows, UNIX, Linux, and NetWare) The user who owns the file is no longer present in the system database.</li> </ul> |

| Condition                        | Description                                                          | Method for filtering the files                                                                                                                                                                                                     | Condition Modifiers                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------------|----------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Attributes</b><br>(continued) |                                                                      |                                                                                                                                                                                                                                    | <ul style="list-style-type: none"> <li>Block Device: (UNIX/Linux only) A disk device file (not a real file).</li> <li>Executable: A file that you can run or execute.</li> <li>Sticky: A special type of executable file that is read into memory when it is first run and never released. Use Scans to find these types of files and set them to "unsticky".</li> <li>SETGID: (UNIX/Linux only, for executable files) When a file is run, its group changes to the group of the file. For example, if a user belongs to User group and runs a file from a different group (e.g., Root), then the user's group will be changed to Root while running the file (for security purposes).</li> <li>SETUID: (UNIX/Linux only, for executable files) When a file is run, the user's effective user will be changed when the file is run. For example, if a file is SETUID =Root, a users effective user ID becomes Root when running the file. This is considered a security hole by UNIX Administrators, so you can define a Profile to search for files that have a SETUID = Root.</li> </ul> |
| <b>Size</b>                      | Use this condition to limit files based on their physical file size. | <ul style="list-style-type: none"> <li>between</li> <li>not between</li> <li>&gt;=</li> <li>&lt;</li> <li>&gt;</li> <li>&lt;=</li> </ul>                                                                                           | <ul style="list-style-type: none"> <li>KB (kilobytes)</li> <li>MB (megabytes)</li> <li>GB (gigabytes)</li> <li>TB (terabytes)</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Type</b>                      | Use this condition to limit files based on pre-defined file types.   | <ul style="list-style-type: none"> <li>in</li> <li>not in</li> <li>=</li> <li>&lt;</li> <li>&gt;</li> </ul>                                                                                                                        | <ul style="list-style-type: none"> <li>Link</li> <li>Directory</li> <li>Fifo</li> <li>Socket</li> <li>Device</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Owner</b>                     | Use this condition to limit files based on their owner.              | <ul style="list-style-type: none"> <li>matches any of</li> <li>matches none of</li> <li>between</li> <li>not between</li> <li>matches</li> <li>does not match</li> <li>&gt;=</li> <li>&lt;</li> <li>&gt;</li> <li>&lt;=</li> </ul> | Enter the name of the owner upon which you want the condition to be based in the <b>Entry</b> field.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

| Condition | Description                                                                        | Method for filtering the files                                                                                                                       | Condition Modifiers                                                                                                                                                                                                                               |
|-----------|------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Length    | Use this condition to limit files based on the logical size of the files contents. | <ul style="list-style-type: none"> <li>• between</li> <li>• not between</li> <li>• &gt;=</li> <li>• &lt;</li> <li>• &gt;</li> <li>• &lt;=</li> </ul> | Enter the file length range. You can qualify both values by selecting from one of the following: <ul style="list-style-type: none"> <li>• KB (kilobytes)</li> <li>• MB (megabytes)</li> <li>• GB (gigabytes)</li> <li>• TB (terabytes)</li> </ul> |

*Associating conditions with groups:*

Associate multiple conditions in a group or groups.

Groups can further define your file filter by including or excluding certain files based on specific criteria.

1. Create a new group using one of the following options.
  - In the **File Filters** tab, right-click a condition and select **New Group** from the pop-up menu.  
The **Create Group** dialog box is displayed.
  - Create each file condition for your file filter.  
As you create a new file condition, the **Create Group** dialog box is displayed.
2. Select the grouping for the file conditions you created. The grouped condition is displayed in the **File Filters** tab.

Select one of the following groups.

| Group             | Description                                                                                                                                  |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>All of</b>     | All inclusive. All the conditions must be met.<br><br>For example: <ul style="list-style-type: none"> <li>• a and b and c</li> </ul>         |
| <b>Any of</b>     | Any one or more of the following conditions will be met.<br><br>For example: <ul style="list-style-type: none"> <li>• a or b or c</li> </ul> |
| <b>None of</b>    | None of the conditions are met.<br><br>For example: <ul style="list-style-type: none"> <li>• not (a or b or c)</li> </ul>                    |
| <b>Not all of</b> | All of the following conditions are not met.<br><br>For example: <ul style="list-style-type: none"> <li>• not (a and b and c)</li> </ul>     |

For example, say you want to create a file filter with the following conditions:

- Name in (\*.gif, \*.jpg)
- Size > 1 MB
- Attribute is Read Only



How you group the conditions produces different results. The following table shows the different results depending on the group you select:

| File                   | Group        | Result  |
|------------------------|--------------|---------|
| A.gif 500KB read only  | • All of     | • False |
|                        | • Any of     | • True  |
|                        | • None of    | • False |
|                        | • Not all of | • True  |
| B.jpg 2 MB read only   | • All of     | • True  |
|                        | • Any of     | • True  |
|                        | • None of    | • False |
|                        | • Not all of | • False |
| C.gif 2 MB read/write  | • All of     | • False |
|                        | • Any of     | • True  |
|                        | • None of    | • False |
|                        | • Not all of | • True  |
| D.txt 500KB read/write | • All of     | • False |
|                        | • Any of     | • False |
|                        | • None of    | • True  |
|                        | • Not all of | • True  |

*Using wildcards in file filters:*

Use wildcard characters within the conditions of file filters to perform pattern matching of file names. A wildcard is a special symbol that represents one or more characters.

You can use wildcards in File Filter conditions to identify multiple files and directories with a single specification. The following table describes the wildcards that are available in Data Manager:

| Wildcard Character | Description                                                                                                                                                            | Examples                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| *                  | Use the asterisk (*) to represent unknown characters in file names. It can act as a wildcard for zero or more characters of any value within a directory or file name. | <ul style="list-style-type: none"> <li>• <b>s*</b> refers to all files that begin with <b>s</b></li> <li>• <b>*.exe</b> refers to all files that have <b>.exe</b> as an extension</li> <li>• <b>/user/*</b> refers to all files under the <b>/user</b> directory</li> </ul>                                                                                                               |
| %                  | Use the percent (%) character to specify a wildcard for a directories within the path of a file. It can act as a wildcard for zero or more directories within a path.  | <ul style="list-style-type: none"> <li>• <b>/home/%/*cmd</b> will match all files with the extension <b>.cmd</b> and located in any of the subdirectories under the <b>/home</b> directory</li> <li>• <b>/usr/%/*</b> will match all files and subdirectories under the <b>/usr</b> directory (for example, <b>/usr/files/c</b>, <b>/usr/c</b>, <b>/usr/files/documents/c</b>)</li> </ul> |

| Wildcard Character | Description                                                                                                                                                    | Examples                                                                                                                                                                                                                                                                                                         |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ?                  | Use the question mark (?) to represent a character within a file name or directory name. It can as a wildcard for a character within a directory or file name. | <ul style="list-style-type: none"> <li>• <code>/home/usr?/</code> will match directories like: <code>/home/usr1/</code>, <code>/home/usr2/</code>, etc.</li> <li>• <code>/usr/temp/temp??.tmp</code> will match files like <code>/usr/temp/temp01.tmp</code>, <code>/usr/temp/temp02.tmp</code>, etc.</li> </ul> |

Data Manager supports standard UNIX shell path name expansion wildcarding within file filter conditions. See the appropriate UNIX man pages for more information about the wildcards that you can use.

## Determining the availability of storage resources (pings)


Use pings to collect information about the availability of the storage resources in your environment. You can view the information collected by pings in Availability and System reports.

Pings enable you to monitor and report on the availability of your storage from a network point of view or from a computer uptime perspective. See the percentage of off-network time due to network problems or system downtime. Define any number of pings that each check the availability of a unique storage resource.

### Creating pings

Create pings to collect information about the availability of the storage resources in your environment. You can define any number of pings to gather information about different storage resources at different times.

To create a ping, complete the following steps:

1. Expand **Data Manager > Monitoring > Pings**.
2. Right-click **Pings** A menu is displayed.
3. Select **Create Ping**.
4. Select the computers whose availability you want to check. "Select computers to ping" on page 365.
5. Schedule when you want a ping to run. "When to run the ping" on page 365.
6. Specify how often you want the server to write out the ping results to the database repository for viewing in reports. "Specify ping options" on page 366.
7. Click  to save the ping. You must save the ping before Tivoli Storage Productivity Center can collect data about storage resources.
8. Enter a name for the ping and click **OK**. An entry for the ping is displayed under the **Pings** node in the navigation tree. The ping is run at the date and time specified in the schedule definition.
9. Use the Job Management panel to manage view information about the runs and jobs that are related to the ping. For example, use the Job Management panel to view the status of a run for a schedule, the date and time when it was run, and the log files of the jobs associated with the run.

**Related reference:**

“Managing schedules, jobs, and runs” on page 295

Use the Job Management panel to view and manage all the schedules, runs, and jobs related to the storage devices that are monitored by IBM Tivoli Storage Productivity Center.

**Select computers to ping:**


Select the computers and clusters that you want to ping.

Use the **Computers** page to select the computers and clusters whose availability statistics you want to collect. The **Available** list box shows the computers, computer groups, or clusters that are currently being monitored and available for inclusion in the ping. After you include storage resources in the ping it is removed from this list and moved to the **Current Selections** list box. The **Current Selections** list box shows the computers, computer groups, or clusters that have been selected for the ping.


**Note:** A cluster name is displayed next to individual computers that are members of a PowerHA or Microsoft Cluster Server cluster. Cluster resource groups are identified with a special icon. Clusters displayed in the **Available** list box cannot be expanded to show member nodes. You can move a cluster to the **Current Selections** list box to view and, if necessary, exclude member nodes.

You can include individual computers, computer groups, or clusters against which to run a ping. See “Working with groups of resources” on page 371 for information about groups.

**• To include storage resources in a ping:**

1. In the **Available** list box, highlight the computers, computer groups, and clusters that you want to include in the ping.
2. Click . The storage resource that you selected is moved to the **Current Selections** list box.

**• To remove storage resources from a ping:**

- Highlight the computers, computer groups, and clusters that you want to remove and click . The storage resource is removed from the ping and moved from the **Current Selections** list box to the **Available** list box.

**• To exclude one or more computers from a selected computer group or cluster:**

1. Double-click the group or name of the cluster in the **Current Selections** list box.
2. Right click the computers that you want to exclude. A menu is displayed.
3. Select **Exclude**. An **Excludes** section is created in the **Current Selections** list box that displays the names of the computers that you have excluded from the ping.

**When to run the ping:**

Schedule when and how often to run a ping.

You have the following options for determining when to run a ping:

- Run the job immediately
- Run the job once at a specified time and date. Use the list boxes to specify the month, day, year, and the hour, minute, and AM or PM.

- Run the job repeatedly according to a schedule you specify. You can do the following:
  - Use the list boxes to specify the month, day, year, and the hour, minute, and AM or PM when the job should begin running.
  - Indicate how often the job should run. Use the list boxes to specify the number of minutes, hours, days, weeks, or months.
  - Select which days you wish to run the job.

Use the fields in the **How to handle timezones** section to indicate the time zone that Data Manager should use when scheduling and running the ping. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:

- 9:00 AM of the time zone where the server resides
- 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

**Use the timezone that the server runs in.** Select this option to use the time zone of the location where the Data server resides.

**Use this timezone.** Select this option to indicate a specific time zone for when to run the job. When you select this option, the list box to the right of the field will activate.

#### **Specify ping options:**

Determine how often the results of pings are written from memory and stored in the database repository.

For greater efficiency, the server temporarily stores ping results in memory. The unit of time you specify in the **Options** tab determines how often these results are externalized to the repository.

You can specify a unit of time in a number of pings, minutes, hours or days.

#### **Define alert values for a ping:**

Set an alert that will be triggered if the ping job fails.

Use the Alert page to define an alert to notify you if a computer that a ping is trying to check is unreachable for a certain amount of time. You can choose to be notified by SNMP trap, IBM Tivoli Enterprise Console event, login notification, Windows event log, UNIX Syslog, running a script, and email. If you select SNMP trap, Tivoli Enterprise Console event, or email, you must also define a destination (for example, an email address). Define this destination by expanding **Administrative Services > Configuration > Alert Dispositions**.

To define an alert for a ping, complete the following steps:

1. Click the **Alerts** tab; the Alert page opens.
2. Define the conditions that will trigger the alert. The default setting is **Computer Is Unreachable More Than**. No other alerting conditions are currently available. Enter the **Value** and **Value Units**. These values specify how many pings occur before an error is generated. By default, **Value Units** is set to **Times**.

3. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition. You can choose from the following actions:

**SNMP trap**

An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

**TEC / OMNIbus Event**

A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, and so on.

**Login notification**

The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

**Login ID**

The ID of the user who receives alert notification upon logging in to the system.

**Windows Event Log, UNIX Syslog**

The ID of the user who receives alert notification upon logging in to the system.

**Event Type (Windows Event Log only)**

The type of event that is recorded to the OS log.

**Facility (UNIX Syslog only)**

The location where the UNIX Syslog is sent. You can select User or Local.

**Level (UNIX Syslog only)**

The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

**Run script**

A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window are passed to the script when the alert is triggered.

**E-mail** An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit e-mail** to change the actual email message.


4. Click **File > Save** to save the ping.
5. Enter a name for the ping and click **OK**. An entry for the ping appears under the **Pings** node.

## Editing pings

Change the settings in an existing ping job.

To edit an existing ping, complete the following steps:

1. Expand **Data Manager > Monitoring > Pings**.
2. Click the name of the ping you want to edit. The content pane displays information about that job.

3. Edit the ping definition by changing the values that appear on the following pages:
  - **Computers.** Use this page to edit the computers, computer groups, and clusters whose availability you want to monitor as part of a ping.
  - **When to Run.** Use this page to edit when a ping is run.
  - **Options** Use this page to edit how often the results of pings are written from memory and stored in the database repository.
  - **Alert.** Use this page to edit how long a storage entity is unreachable before you are notified that the ping has failed.
4. Click  to save the updated ping definition.

## Data collection groups

Groups enable you to define a set of related resources such as a computer, fabric, or storage subsystem once and save these selections as a named group for future use in scans, pings, and probes, making data collection more efficient. For example, you can create a group that contains a number of related computers. You can then reference that group within a monitoring job when you want to collect data about all those computers at the same time.

Each type of group is represented by a node under the **Monitoring > Groups** and **IBM Tivoli Storage Productivity Center > Monitoring > Probes** sections. You can expand the tree for each type of group to see what groups exist for that type. Resources can belong to only one group at a time. For example, if a computer already belongs to a computer group and you attempt to add it to a different computer group, the system removes it from the first computer group before adding it to the new computer group (this includes removing it from the default group if it had never been previously assigned to a user defined group). This is to prevent a resource from being counted twice statistically.

Additionally, you can use Data Manager for Chargeback to create invoices for these groups. When implementing your chargeback system, it is recommended that you group the IT resources within your organization into functional groups. Accurately identifying the types and locations of your stored data and resources is the first step in implementing an effective and equitable chargeback policy. Defining groups with Data Manager enables you to classify your data and resources into billable entities, which in turn provides information on who is using what storage. This enables you to bill specific groups according to their usage of storage space.

The following table lists and describes the types of groups you can create.

| Group Type          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Data Manager</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Computer</b>     | <p>A logical grouping of related computers. A computer can belong to one computer group only. A computer will automatically be removed from a group if you add it to another group. You can use computer groups to monitor computers that are related to each other or belong to a common group such as a department, location, territory, and so on.</p> <p><b>Examples of computer groups:</b> Sales, Accounting, Development, IT, West, East, Midwest, NAS Filers, Remote Servers</p> |

| Group Type                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Filesystem                        | <p>A grouping of file systems across multiple computers. A file system can belong to one file system group only. A file system will automatically be removed from a group if you add it to another group.</p> <p><b>Examples of file system groups:</b> C:\, D:\, /boot, /home, /usr</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Directory                         | <p>A grouping of directories across multiple computers. You can include a directory in multiple directory groups.</p> <p><b>Examples of directory groups:</b></p> <p>WindowsTempDir Included Computer Name:Directory -<br/> a:\machine1\winnt\temp, b:\machine2\winnt\temp,<br/> c:\machine3\winnt\temp</p> <p><b>Note:</b> To monitor directories from the IBM Tivoli Storage SAN File Systems, you must create a directory group and add the desired directories as follows:</p> <pre data-bbox="699 743 1276 772">/ &lt;SAN File System cluster name&gt; / &lt;directory name&gt;</pre> <p>where &lt;SAN File System cluster name&gt; represents the name of the SAN File System cluster that contains the directory that you want to add and &lt;directory name&gt; represents the name of the directory you want to add (for example, /IBM/myFilesetName/myDirectoryName, /IBM/myDirectoryName).</p> |
| User                              | <p>A grouping of related users. A user can belong to one user group only. A user will automatically be removed from a group if you add it to another group. Users that were discovered by Tivoli Storage Productivity Center can be included in a user group.</p> <p><b>Note:</b> The user groups that you define in Data Manager for Databases will also be available in Data Manager, and vice versa.</p> <p><b>Examples of user groups:</b> Administrators, Developers, Marketing staff, HR staff</p>                                                                                                                                                                                                                                                                                                                                                                                                  |
| OS User Group                     | <p>A grouping of OS user groups that were defined outside of Tivoli Storage Productivity Center at the operating system level. OS user groups are the operating system groups that Tivoli Storage Productivity Center finds that own files and directories in the environment.</p> <p><b>Examples of OS user groups:</b> Windows, Linux, UNIX</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Data Manager for Databases</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Computer                          | <p>A logical grouping of related computers. A computer can belong to one computer group only. A computer will automatically be removed from a group if you add it to another group. You can use computer groups to monitor computers that are related to each other or belong to a common group such as a department, location, territory, and so on.</p> <p><b>Examples of computer groups:</b> Sales, Accounting, Development, IT, West, East, Midwest, NAS Filers, Remote Servers</p>                                                                                                                                                                                                                                                                                                                                                                                                                  |

| Group Type                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Databases-<br>Tablespaces | A grouping of RDBMS databases or tablespaces. You can use database-tablespace groups to monitor entire databases or tablespaces (depending on the RDBMS) across multiple instances. A database or tablespace will automatically be removed from a group if you add it to another group. A database or tablespace can belong to one database-tablespace group only.<br><br><b>Examples of database-tablespace groups:</b> DB2 databases, system tablespaces                                    |
| Table                     | A grouping of related RDBMS tables.                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| User                      | A grouping of related users. A user can belong to one user group only. A user will automatically be removed from a group if you add it to another group. Users that were discovered by Tivoli Storage Productivity Center can be included in a user group.<br><b>Note:</b> The user groups that you define in Data Manager for Databases will also be available in Data Manager, and vice versa.<br><br><b>Examples of user groups:</b> Administrators, Developers, Marketing staff, HR staff |
| <b>Disk Manager</b>       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Storage Subsystem         | A storage subsystem group can contain subsystems and SAN Volume Controllers.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Fabric Manager</b>     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Fabric                    | A grouping of fabrics.                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Tape Manager</b>       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Tape Libraries            | A grouping of tape libraries.                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## Default groups

Default groups for each storage resource are also provided within each group type node. As IBM Tivoli Storage Productivity Center discovers these resources within your enterprise through discovery, probes, and scans, the resources are automatically added to the corresponding default groups.

The following table lists the names of the default groups for each group type.

| Manager                         | Group Type                | Default Group Name               |
|---------------------------------|---------------------------|----------------------------------|
| <b>Data Manager</b>             | Computer                  | TPCUser.Default Computer Group   |
|                                 | Filesystem                | TPCUser.Default FS Group         |
|                                 | Directory                 | n/a                              |
|                                 | User                      | TPCUser.Default User Group       |
|                                 | OS User Group             | TPCUser.Default OSGroup Group    |
| <b>Data Manager - Databases</b> | Computer                  | TPCUser.Default Computer Group   |
|                                 | Databases-<br>Tablespaces | TPCUser.Default Tablespace Group |
|                                 | Table                     | n/a                              |
|                                 | User                      | TPCUser.Default User Group       |



| Manager        | Group Type         | Default Group Name                      |
|----------------|--------------------|-----------------------------------------|
| Disk Manager   | Storage Subsystems | TPCUser.Default Storage Subsystem Group |
| Fabric Manager | Fabrics            | TPCUser.Default Fabric Group            |
| Tape Manager   | Tape Library       | TPCUser.Default Tape Group              |

## Working with groups of resources

Create resource groups to logically arrange resources and to permit actions against the entire group.



### Creating a group

You can create named groups that represent a set of resources, such as computers, databases, table spaces, tables, and users, to later use in pings, scans, and probes.

To create a new group, perform the following steps.

1. Expand **Monitoring > Groups**. For example, **Data Manager > Monitoring > Groups**.
2. Right-click the type of group you want to create and select **Create group type**, where *group type* represents the group you want to create, such as a **Create Computer Group**, **Create Filesystem Group**, or **Create Table Group**.
3. In the *Create group type* Groups page, select the resources in the **Available** column that you want to include in the group. Hold down the Ctrl key as you click to select multiple resources.



**Note:** For computer and file system groups in Data Manager, a **Group By Cluster** check box is displayed. When this option is not selected, the **Available** pane displays a cluster name next to any clustered computers and cluster resource groups. When this option is selected, any computer or cluster resource group that belongs to a PowerHA or MSCS cluster is grouped under its cluster name. Cluster resource groups are identified with a special icon along with the name of the current host node. All other computers will be grouped under **Non-Clustered Computers**.

4. Click  to move the selected resources to the **Current Selections** column on the right. Click  to move resources from the **Current Selections** column and move them to the **Available** column.
5. Optional: Type a description of the group in the **Description** field.
6. Click the save icon in the menu bar to save the group.
7. Type a descriptive name for the group and click **OK**.

### Creating tape library groups

Create groups of tape libraries that have similar attributes to make it easier to perform storage management tasks on the grouped libraries.

1. Expand the **Tape Manager > Monitoring > Groups** node and right-click **Tape Library**.
2. Select **Create Tape Library Groups**.
3. In the *Create Tape Library Groups* page, select the libraries in the **Available** column to include the libraries in the group.  
Hold down **Ctrl** to click multiple libraries.

4. Click  to move selected libraries to the **Current Selections** column on the right.
5. Optional: Enter a description of the tape library group in the **Description** field.
6. Click  to type a descriptive name for the group in the **Save As** dialog. Then click **OK** to save the group.


The new tape library group is listed under the **Tape Library** node in the navigation tree.

## Grouping fabrics

You can group fabrics to run reporting jobs against all the fabrics in the group. A fabric can belong to multiple groups at a time.

For more information on groups, see “Data collection groups” on page 368.



To create a fabric group, complete the following steps:

1. Expand **Fabric Manager > Monitoring > Groups**, right-click **Fabric**, and select **Create Fabric Group**.
2. Enter a description for the group, and from the **Available** column, select the fabrics to be added to the group. Move them to the **Current Selections** column by clicking .
3. To save the group, click **File > Save**.
4. In the **Save As** window, enter the group name and click **OK**.

## Editing groups

When you update the resources that are included within a group, those updates are automatically picked up by the jobs that reference that group.

To edit an existing group, perform the following steps.

1. Expand **Monitoring > Groups**. For example, to edit groups within Disk Manager, expand **Disk Manager > Monitoring > Groups > Storage Subsystem**.
2. Expand a group type to view the list of groups defined for that type. For example, to view storage subsystem groups, expand the **Storage Subsystem** node under Disk Manager.
3. Click the name of the group you want to edit. Details for the selected group are displayed in the Create *group type* Groups page, where group type is the type of group (for example, Directory or Computer).
4. Edit the values defined for the group as desired. Click  to move the selected resources to the **Current Selections** column on the right. Click  to move resources from the **Current Selections** column and move them to the **Available** column.
5. Click the save icon in the menu bar to save the group.

## Creating directory groups

Create and use directory groups within scans to monitor and gather information about specific directories within filesystems across a number of machines.

**Note:** IBM Tivoli Storage SAN File Systems information: To monitor directories from the SAN File System, you must create a directory group and add the desired directories as follows:

*/<IBM Tivoli Storage SAN File Systems cluster name>/<directory name>*

where:

- *<IBM Tivoli Storage SAN File Systems cluster name>* represents the name of the IBM Tivoli Storage SAN File Systems cluster that contains the directory you want to add
- *<directory name>* represents the name of the directory you want to add

Examples:

- /IBM/myFilesetName/myDirectoryName
- /IBM/myFilesetName
- /IBM/myDirectoryName

To create a directory group, perform the following steps.

1. Expand **Data Manager > Monitoring > Groups > Directory** or click **Create Directory Group** from the **Directory Groups** tab of the Scan window. The Create Directory Group window opens.
2. Use the New Directory window to edit or add directories to a directory group.

| Field                              | Description                                                                                                                                                                                                                                                                     |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Creator</b>                     | Displays the creator of the group.                                                                                                                                                                                                                                              |
| <b>Name</b>                        | Displays the name of the group you are adding or editing. This name can be helpful in determining the content of the group. For example, a directory group named UnixUserHomes might consist of the /home directory on all the UNIX and Linux computers within an organization. |
| <b>Description</b>                 | A description of the group you are adding or editing.                                                                                                                                                                                                                           |
| <b>New Computer, New Directory</b> | Click these buttons to add a directory or computer to the group.                                                                                                                                                                                                                |

#### Viewing directory groups:

When adding directories to a group, you can determine how you want to view groups of directories. You can choose either of the following options:

- **Directories by computer.** Use this option to display the directories in the directory group according to the computer on which they exist.
  1. Expand Data Manager.
  2. Expand Monitoring.
  3. Expand Groups.
  4. Expand Directory.
  5. Double-click a directory object.
  6. In the right pane, under View, select Directories by Computer which is the default setting.

**Note:** You might need to try more than one directory object to find one with data.

- **Computers by directory.** Use this option to display the computers in the group according to the directories included in the directory group.

1. Expand Data Manager.
2. Expand Monitoring.
3. Expand Groups.
4. Expand Directory.
5. Double-click a directory object.
6. In the right pane, under View, select Computers by Directory.

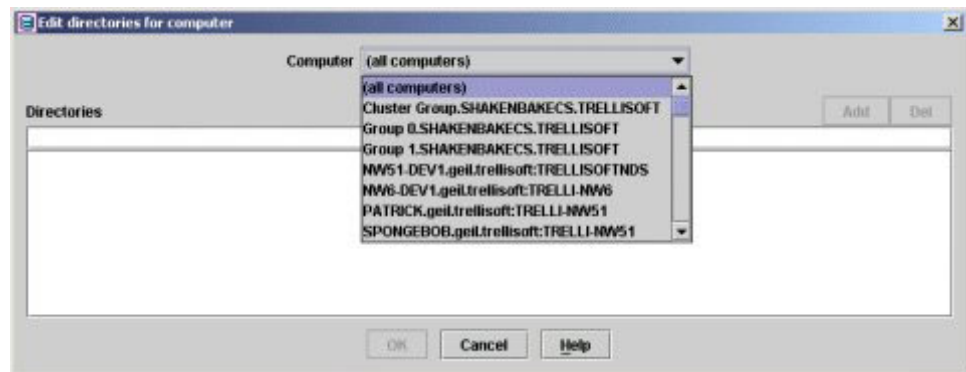
**Note:** You might need to try more than one directory object to find one with data.

**Tip:** Choosing these different buttons enables you to add directories to a directory group in different ways. For example, if you want to monitor or receive alerts on different directories on multiple computers, choose the Directories by Computer option. If you want to monitor or receive alerts on the same directory across multiple computers, choose the Computers by Directory option.

#### Adding new directories by computer:

If you want to add new directories to a directory group by first selecting the computers on which they exist, perform the following steps:

1. Select **directories by computer** on the Create Directory Group window.
2. Click **New computer**. The Edit Directories for Computer window opens. This window enables you to select the computer for which you want to monitor specific directories, and enter the directories for the selected computer.




3. Click the **Computer** list box.
4. Select the computer for which you want to monitor specific directories.
5. Enter the directory you want to add to the group in the **Directories** field. For example: c:\or /home/\*.
6. Click **Add**. The directory is added to the list box.
7. Click **OK** when you are done adding directories. The New Directory Group window opens. The computer you added will be displayed as a node in the display area, with the specified directories under the node.

#### Adding new computers by directory:

If you want to add new directories to a directory group by first specifying the directory and then selecting the computers on which they exist, perform the following steps:

1. Select **computers by directory** on the Create Directory Group window.
2. Click **New computer**. The Edit Computers for Directory window opens. This window enables you to enter the directories you want to add to the directory group and specify which computers contain those directories.

3. Enter the directory you want to add to the group in the **Directory** field. For example: c:\, /home/\*, \user\\*, \windows\temp\, /usr/\*.
4. Select the computers you want to include in the group and click . The computer will be displayed in the **Current Selections** column. Select the **All computers** option to include all computers that contain the directory in the Directory Group.
5. Click **OK**. The New Directory Group window opens. The directory you added will be displayed as a node in the display area, with the selected computers under the node.

---

## Collecting performance data

Use the tasks in this section to collect performance data about your storage subsystems, fabric switches, and tape libraries.

### Monitoring switch performance

You can monitor the performance of those switches in your fabrics from which you can collect performance data.

#### Creating a switch performance monitor

You can monitor the performance of selected fabric switches. Fabric Manager can collect performance data for the switches and generate reports from the results of those monitors.

Report data is collected on data rates, operations rates, switch port errors, and other performance statistics. For information about generating those reports, see “Creating switch performance reports” on page 974.

**Note:** A fabric probe is now required after a CIMOM discovery before a switch performance monitor job can be run.


To create a switch performance monitor, complete the following steps:

1. Expand **Fabric Manager > Monitoring**, right-click **Switch Performance Monitors**, and select **Create Switch Performance Monitors**.
2. In the **Switches** tab, enter a description for the performance monitor, and from the **Available** column, select the switches to be added to the monitor definition.

Move them to the **Current Selections** column by clicking .

**Note:** Only those switches from which Fabric Manager can collect performance data are listed in the **Available** column. If a switch is added to a performance monitor, it is no longer listed as an available switch.

3. In the **Sampling and Scheduling** tab, specify an interval for the performance averages, the duration of the data collections, and a schedule for the monitor. You can begin the monitor immediately, schedule it to begin at a later time, and specify whether to run the monitor repeatedly.
4. In the **Alert** tab, specify the action to be taken if the monitor fails. See “Running scripts with alerts” on page 507 for a description of these triggered actions.

5. Click  to save the performance monitor. You must save the monitor before Tivoli Storage Productivity Center can collect data about storage resources.

6. Enter a name for the performance monitor and click **OK**. An entry for the monitor is displayed under the **Switch Performance Monitors** node in the navigation tree. The monitor is run at the date and time specified in the schedule definition.
7. Use the Job Management panel to manage view information about the runs and jobs that are related to the performance monitor. For example, use the Job Management panel to view the status of a run for a schedule, the date and time when it was run, and the log files of the jobs associated with the run.

### Editing a switch performance monitor

Edit a saved switch performance monitor to change settings, such as which switches to include in the collection.

1. Expand **Fabric Manager > Monitoring > Switch Performance Monitors**. There will be a node for each saved monitor definition.
2. Click the name of the monitor you want to edit.
3. Edit the monitor definition.
4. To save the edited monitor, click **Save**.

## Monitoring storage-subsystem performance

Discover, probe, define performance thresholds, and view performance data and alerts to monitor a storage system.

1. Configure the storage system for monitoring by IBM Tivoli Storage Productivity Center.
  - If you want to monitor IBM System Storage DS8000 systems, IBM System Storage SAN Volume Controller, IBM Storwize V7000 Unified, IBM Storwize V7000, or IBM XIV Storage Systems, use the Configure Devices wizard to enter connection information for their native interfaces. See “Configure storage subsystems” on page 306 for details of how to use the Configure Devices wizard.
  - If you want to monitor storage systems that use CIM object manager (CIMOM) connections, there are a number of ways to enter the connection information into Tivoli Storage Productivity Center. If the CIMOM is located in the same subnet as the Device server, it can be discovered through Service Location Protocol (SLP). You can manually add the CIMOM to Tivoli Storage Productivity Center by expanding the **Administrative Services > Data Sources** node and clicking **CIMOM Agents**. See Adding a CIM agent manually for details of how to add the CIMOM. Alternatively, you can also use the Configure Devices wizard to enter the CIMOM information.
2. Discover the storage system.

Discovery gets general information about the devices that affiliated with the CIMOMs. These devices are switches and other storage systems such as IBM System Storage DS5000 and other DS series systems, and IBM TotalStorage Enterprise Storage Server. You can run discovery for CIMOMs by expanding the **Administrative Services > Discovery > CIMOM** node. Alternatively, you can also use the Configure Devices wizard to run discovery for CIMOMs.

The discovery step is not required for the storage systems that are connected to Tivoli Storage Productivity Center using native interfaces.
3. Probe the storage system.

Probes collect information about logical devices, such as volumes and virtual disks, that are associated with the storage system. Probes also collect information about hosts that are mapped to the logical devices.

If you use the Configure Devices wizard, storage subsystems are included automatically in probes. To create probes manually, expand **IBM Tivoli Storage Productivity Center > Monitoring > Probes**. For Data Manager for Databases, expand **Data Manager for Databases > Monitoring > Probes**.


4. Define storage-system performance thresholds.  
Thresholds draw your attention to exceptions to normally expected values in the data that you collect.
5. Collect storage-system performance data.  
Collect performance data for an individual storage system.
6. View storage system performance data.

### Creating storage subsystem performance monitor

Create a subsystem performance monitor to gather performance data for storage systems.

**Note:** If you schedule a performance monitor to collect data about a storage subsystem at the same time you plan to create many volumes on that subsystem, the performance of the volume creation job and the general performance of the associated CIMOM might decrease. Consider scheduling a performance monitoring job at a different time than when you plan to create many volumes on a subsystem.

Complete the following steps to collect performance data for a storage system:

1. In the left pane, expand **Disk Manager → Monitoring**.
2. Right-click **Subsystem Performance Monitor** and click **Create Performance Monitor**.  
The Create Storage Subsystem Performance Monitor window opens in the right pane.
3. Optional: In the **Description** field, type a description of the threshold.
4. Click the **Storage Subsystems** tab. The Storage Subsystems page is displayed and all storage subsystems that have been discovered are listed.
5. In the **Available subsystems** field, click storage subsystems to select them for collection. Then click .

The selected storage subsystems are displayed in the **Current selections** field.

6. Verify that you have created and run a probe for each selected storage subsystem, as described in “Creating probes” on page 336.  
Storage-subsystem reports use data gathered by probes, including data on storage assets in the enterprise, such as computers, storage subsystems, disk controllers, hard disks, and file systems. Data collection fails if a selected system has not yet been probed.
7. Click the **Sampling and Scheduling** tab. The **Sampling and Scheduling** page is displayed.
8. In the **Sampling** area, specify the length of the sampling window and how frequently the performance data is gathered.
  - a. In the **Interval length** field, select the length of the sampling window.  
The interval length is the number of minutes over which performance data is averaged. For example, if you select 15 minutes, all performance-data samples gathered from the storage subsystem represent an average of the performance of the subsystem over the previous 15 minute interval.
  - b. In the **Duration** field, specify how long the monitoring task will run.

You can specify that the task runs indefinitely or specify the length of time that the performance data is gathered.

- c. In the **Scheduling** area, click either **Begin immediately** or **Schedule to begin later**.

If you click **Schedule to begin later**, specify the date and time that you want the data collection to start. You also can specify that the data collection recurs, how frequently it recurs, and whether certain days of the week are excluded.

**Important:** When you specify a start time for using performance monitor data in a plan, a report, or an optimization plan, the start time might not coincide with the start of a performance monitor data collection interval. If the requested start time is after the start of a data collection interval and the duration of the data collection interval is long enough so that it ends at or after the requested start time, Tivoli Storage Productivity Center includes all data from the collection interval that began before the requested start time. This might include data collected before the start time that you specified. This processing ensures that the maximum amount of available performance data is utilized. For example, if a performance monitor collects data every 20 minutes starting at 01:00, but you request performance data starting at 1:15, Tivoli Storage Productivity Center still analyzes performance data starting at 01:00 that has a duration of 20 minutes.

9. Click the **Alert** tab. The **Alert** page is displayed. In the **Condition** area **Monitor Failed** is the triggering condition.
10. In the **Triggered Actions** area, select actions to trigger when the data collection task fails.

You can select from the following check boxes:

- **SNMP Trap**
- **TEC / OMNIBUS Event**
- **Login Notification**
- **Windows Event Log**
- **Run Script**
- **Email**

Depending on what action you select, you might have other choices to make. Suppose that you select the **Windows Event Log** check box, for example. The **Event Type** field becomes activate to specify a severity for the event in the Windows event log.

11. Click **File** → **Save**. You must save the performance monitor before Tivoli Storage Productivity Center can collect data about storage resources.
12. Enter a name for the performance monitor and click **OK**.

An entry for the monitor is displayed under the **Subsystem Performance Monitors** node in the navigation tree. The monitor is run at the date and time specified in the schedule definition.

13. After the monitoring job completes, expand the name of the storage-subsystem performance monitor.

The following information is displayed:

- A status icon that indicates whether the task completed successfully
- Number of times that the task ran
- Date that the task ran
- Time that the task began




14. To view additional information, click the task information.

In the right pane, the following additional information is displayed:

- Time that the task finished
- Number of jobs in the task
- Number of jobs that failed

A table contains information about the storage subsystems in the data collection.

15. In the table, click a storage subsystem, and then click  to display the contents of the log file.

16. Use the Job Management panel to manage view information about the runs and jobs that are related to the monitor. For example, use the Job Management panel to view the status of a run for a schedule, the date and time when it was run, and the log files of the jobs associated with the run.

## Example log file

The following example is a log file for a SAN Volume Controller:

```
Job log file: <Location of log file>
<date><time>HWNPM21131 The performance monitor for device
 <name> (<IP address>:<?>) is starting in an active state
<date><time>HWNPM21151 Monitor Policy:name="svcProbe",
creator="administrator", description=""
<date><time>HWNPM21161 Monitor Policy:retention period:
sample data=14 days,hourly data=30 days,daily data=90 days
<date><time>HWNPM21171 Monitor Policy:interval length=900 secs,
frequency=900 secs, duration=1 hours
<date><time>HWNPM21181 Threshold Policy:name=
"Default Threshold Policy for SVC",creator="System",
description="Current default performance threshold policy for
SAN Volume Controller"
<date><time>HWNPM21191 Threshold Policy:retention period:
exception data=14 days
<date><time>HWNPM21201 Threshold Policy:threshold name=
Total I/O Rate Threshold,enabled=no,
boundaries=-1,-1,-1,-1 ops/sec
<date><time>HWNPM21201 Threshold Policy:threshold name=
Total Data Rate Threshold,enabled=no,
boundaries=-1,-1,-1,-1 MB/sec
<date><time>HWNPM21201 Threshold Policy:threshold name=
Total Backend I/O Rate Threshold,
enabled=no, boundaries=3,2,1,0 ops/sec
<date><time>HWNPM21201 Threshold Policy:threshold name=
Total Backend Data Rate Threshold,
enabled=yes, boundaries=3,2,1,0 MB/sec
<date><time>HWNPM21201 Threshold Policy:threshold name=
Overall Backend Response Time Threshold,
enabled=yes,boundaries=-1,-1,-1,-1 ms/sec
<date><time>HWNPM02001 This operation (startPerfCollection())
on Performance Manager was successful.
<date><time>HWNPM22001 Successfully retrieved the configuration
data for the storage system.
Found 1 I/O Groups,2MDisk Groups,25MDisks, and 1 VDisk
<date><time>HWNPM21231 Performance data was collected and processed
successfully.
27 performance data records were inserted into the database.
<date><time>HWNPM21231 Performance data was collected and processed
successfully.
54 performance data records were inserted into the database.
<date><time>HWNPM21261 The performance monitor for device
<name> (<IP address>:<?>) is stopping because its
intended duration has elapsed.
```

Descriptions of variables in the log file include:

- <date> is the date, for example, 1/23/06.
- <time> is the time, for example, 12:45:20 PM.
- <name> is the device ID of the storage subsystem, for example, SVC-2145-RedC3-IBM.
- <IP address> is the IP address of the storage subsystem, for example, 9.47.97.159.

This log also provides the following information:

- The monitoring task was targeted against a single SAN Volume Controller, SVC-2145-RedC3-IBM.
- The name of the monitoring task is **svcProbe**, and it was created by a user with an ID **administrator**.

- The retention period for exception data.
- The user had provided the following values for the task: interval length: sampling window of 900 seconds (15 minutes) and duration of one hour.
- The monitoring task successfully retrieved the configuration data from the SAN Volume Controller. The task associated the following logical entities with the SAN Volume Controller:
  - One I/O group
  - Two managed-disk groups (MDisk groups)
  - 25 managed disks (MDisks)
  - One virtual disk (VDisk)
- The monitoring task successfully collected and processed data for two performance metrics. A total of 81 performance records were written to the database.
- The monitoring task stopped when its duration (one hour) had elapsed.

### Editing a storage-subsystem performance monitor

Edit a storage-subsystem performance monitor.

Complete the following steps to edit storage subsystem performance monitor:

1. In the left pane, expand Disk Manager > **Monitoring** > **Subsystem Performance Monitors**. The list of existing subsystem performance monitors is displayed.
2. Click the monitor that you want to edit. The Edit Subsystem Performance Monitor window opens.  
The Create Storage Subsystem Performance Monitor window opens in the right pane.
3. Make any necessary changes.
4. Click **File** → **Save**.

### Monitoring tape libraries

Monitor the performance of tape libraries to be aware of events that occur, and to respond to alerts.

1. Create one or more tape library groups, as described in “Creating tape library groups” on page 371.
2. Create a probe job to collect data about your tape libraries and library groups, as described in “Creating tape library probes” on page 338.
3. View alerts that are generated by your tape libraries, as described in “Viewing alert logs” on page 401.

### Viewing performance data

View collected performance data and recognized threshold exceptions in various user-friendly formats to stay aware of usage trends and exceptions.

Display data in a tabular format (actual values of the collected metrics), a graphical format (primarily as line-graphs), or export the data to a file. Select whether recent data, historical data, or information about threshold exceptions (constraint violations) is displayed.

Select a device or a device component, select a particular metric, and specify a time range to plot the data over the time range.

Use constraint violation reports to view a list of any threshold exceptions that occurred during a particular time range. Bar graphs of exceptions are commonly used to show the total number of threshold violations.

## Viewing storage-subsystem performance data

View collected storage-subsystem performance data and recognized threshold exceptions to stay aware of storage-subsystem usage trends and exceptions.

1. In the left pane, expand **Disk Manager** → **Reporting** → **Storage Subsystem Performance**.

A list of reports is displayed under **Storage Subsystem Performance node**.

The following reports generate information concerning SAN Volume Controller models:

- By Storage Subsystem
- By Controller
- By I/O Group
- By Module/Node
- By Array
- By Managed Disk Group
- By Volume
- By Managed Disk
- By Port
- Constraint Violations

2. Click **By I/O Group** to view information about I/O groups.



In the right pane, the Selection page is displayed. By default, all applicable performance metrics are selected for inclusion in the report.

3. Select the data or performance metrics that you want to exclude from the report in the **Included Columns** field. See “Performance metrics” on page 1186 for a list of performance metrics.

- a. Click the items that you want to exclude from the report.

- b. Click .

Excluded items are displayed in the **Available Columns** field.

- c. Select a metric and click either  or  to change the order of the selected metric in the report.

4. Type the number of I/O groups to include in the report, in the **Return maximum of** field.

Limit the number of I/O groups to reduce the amount of performance information in the report. By default, this value is set to 2500.


5. Click **Selection** to open the Select Resources window.
6. Select check boxes to select I/O groups for the report, then click **OK**.
7. Click **Filter** to open the Edit Filter window.
8. Specify filter conditions for the report.
9. Click **Generate Report**.

An I/O Group page opens and the report is displayed.

10. Optional: View the report in chart format.

- a. Click , to specify charting options, then click **OK**.

b. Click **Customize this chart** to further customize the chart.

11. Optional: Click  and type a name to save the report settings for future use.

The report is saved under the **IBM Tivoli Storage Productivity Center** node.

## Viewing tape library information

View collected tape library data to stay aware of discovered tape libraries and usage information.

- Click **Tape Manager**. Left-click **Tape Libraries**. In the right pane, the Tape Libraries page is displayed, which lists the discovered tape libraries.  
Select a library from the table and then click one of the buttons to see further details about that library. For example, select a library and click **Drives** to open a page that lists all the drives for that library.  
You can also right-click a library in the **Tape Libraries** table and select **View Details**.
- Use the **View Details** feature in the topology viewer, as described in “Working with tape libraries in the topology viewer” on page 554.

---

## Working with alerts

Create an alert, configure how the alert is publicized, or edit, delete, disable, rename, or view an alert.

### Related reference:

“Performance thresholds” on page 1176

Performance thresholds are triggering conditions which are used to monitor a component with user-defined values.

## Creating alerts

You can define alerts so that IBM Tivoli Storage Productivity Center notifies you when a specified event occurs. The specific triggering conditions that you can use depend on the type of storage resource that you are monitoring.

**Note:** Alerts are not generated in a Tivoli Storage Productivity Center instance for actions that you perform from that instance. For example, if you start Tivoli Storage Productivity Center and use Disk Manager to assign or unassign volumes for a subsystem, you will not receive alerts for those volume changes in that instance. However, if you assign and unassign volumes outside of that Tivoli Storage Productivity Center instance, an alert is generated.

### Related reference:

“Performance thresholds” on page 1176

Performance thresholds are triggering conditions which are used to monitor a component with user-defined values.

## Creating computer alerts


You can use these alerts to monitor computers for changes in configuration or status. When a condition is detected during a probe, an alert is generated. Computer alerts can be applied to computers, computer groups, and clusters.

If necessary, complete the following prerequisites before you begin the task:

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.

- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services > Configuration > Alert Disposition**.

To create a computer alert, complete the following steps:

1. In the Navigation Tree pane, expand **Data Manager > Alerting**.
2. Right-click **Computer Alert** and click **Create Alert**. The Create Alert window opens and the Alert page is displayed.
3. Select the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering-Condition** area, specify the event that will trigger the alert:
  - a. In the **Condition** field, select the triggering condition.
  - b. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be unavailable.
  - c. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be unavailable.
6. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see “Triggered actions for alerts” on page 1172.
7. Click the **Computer** tab; the Computers page opens.
8. Select the storage entities to which you want the alert to apply:
  - a. In the **Available** field, click the storage entities.
  - b. Click . The storage entities are moved to the **Current Selections** field.
9. Save the alert:
  - a. Click **File > Save**.
  - b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

## Creating directory alerts


You can use these alerts to be notified when either a directory is not found or a directory consumes more than a specified amount of storage space. When a condition is detected during a scan, an alert is generated. These alerts can be applied to directories only.

Before you can create a directory alert, you must have completed the following prerequisites:

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services > Configuration > Alert Disposition**.
- Directories must be grouped into a directory group before they can be included in an alert.

To create a directory alert, complete the following steps:

1. In the Navigation Tree pane, expand **Data Manager > Alerting**

2. Right-click **Directory Alert** and click **Create Alert**. The Create Alert window opens and the Alert page is displayed.
3. Select the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering-Condition** area, specify the event that will trigger the alert:
  - a. In the **Condition** field, select the triggering condition.
  - b. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be unavailable.
  - c. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be unavailable.
6. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see “Triggered actions for alerts” on page 1172.
7. Click the **Directories** tab; the Directories windows opens.
8. Select the storage entities to which you want the alert to apply:
  - a. In the **Available** field, click the storage entities.
  - b. Click . The storage entities are moved to the **Current Selections** field.
9. Save the alert:
  - a. Click **File > Save**.
  - b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

## Creating file system alerts


You can use these alerts to monitor file systems for available free space, changes in configuration, or a file system being removed or unmounted. When the condition is detected during a scan, a probe, or upon an SMNP event, an alert is generated.

If necessary, complete the following prerequisites before you begin the task:

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services > Configuration > Alert Disposition**.

File system alerts can be applied to file systems, file system groups, computers, computer groups, storage systems, storage system groups, and clusters. To create a file system alert, complete the following steps:

1. In the navigation tree, expand **Data Manager > Alerting**.
2. Right-click **Filesystem Alert**, and then click **Create Alert**. The Create Alert panel opens and the Alert page is displayed.
3. Select the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering-Condition** area, specify the event that will trigger the alert:
  - a. In the **Condition** field, select the triggering condition.
  - b. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be unavailable.

- c. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be unavailable.
6. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see “Triggered actions for alerts” on page 1172.
7. Click the **Filesystems** tab; the Filesystems window opens.
8. Select the storage entities to which you want the alert to apply:
  - a. In the **Available** field, click the storage entities.
  - b. Click . The storage entities are moved to the **Current Selections** field.
9. Save the alert:
  - a. Click **File > Save**.
  - b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

### Creating other NAS alerts

You can use other network-attached storage (NAS) alerts to monitor systems that include NAS and are not included in other alert types. When a change in configuration or status is detected during a probe or from an SNMP event, an alert is generated.


If necessary, complete the following prerequisites before you begin the task:

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services > Configuration > Alert Disposition**.

Other NAS alerts can be applied to NAS clusters, nodes, storage systems, and storage system groups. To create a NAS alert, complete the following steps:

1. In the navigation tree, expand **Data Manager > Alerting**.
2. Right-click **Other NAS Alerts**, and then click **Create Alert**. The Create Alert panel opens and the Alert page is displayed.
3. Select the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering-Condition** area, specify the event that will trigger the alert:
  - a. In the **Condition** field, select the triggering condition.
  - b. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be unavailable.
  - c. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be unavailable.
6. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see “Triggered actions for alerts” on page 1172.
7. Click the **NAS** tab; the NAS page opens.
8. Select the storage entities to which you want the alert to apply:
  - a. In the **Available** field, click the storage entities.



- b. Click . The storage entities are moved to the **Current Selections** field.
9. Save the alert:
  - a. Click **File > Save**.
  - b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.


## Creating instance alerts

You can use these alerts to monitor instances for changes. When the condition is detected during a probe, an alert is generated. These alerts can be applied to instances, computers, and computer groups.

If necessary, complete the following prerequisites before you begin the task:

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services > Configuration > Alert Disposition**.

To create an instance alert, complete the following steps:

1. From the Navigation Tree pane, expand **Data Manager for Databases > Alerting**.
2. Right-click **Instance Alerts** and click **Create Alert..** The Create Alert window opens and the Alert page is displayed.
3. Select the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** area, specify the event that will trigger the alert:
  - a. In the **Rdbms Type** field, select the type of database application. You can choose **All Rdbms**, **Oracle**, **SQL\Server**, **Sybase**, or **UDB**.
  - b. In the **Condition** field, select the triggering condition.
  - c. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be unavailable.
  - d. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be unavailable.
6. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see “Triggered actions for alerts” on page 1172.
7. Click the **Instances** tab; the **Instances** window opens.
8. Select the storage entities to which you want the alert to apply:
  - a. In the **Available** field, click the storage entities.
  - b. Click . The storage entities are moved to the **Current Selections** field.
9. Save the alert:
  - a. Click **File > Save**.
  - b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

## Creating database and tablespace alerts

You can use these alerts to monitor database and tablespaces for changes. When a scans or probes are run, events are recognized and alerts triggered. These alerts can be applied to databases, database groups, table spaces, table space groups, instances, computers, and computer groups.

If necessary, complete the following prerequisites before you begin the task:

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services > Configuration > Alert Disposition**.

**Note:** (DB2 only) Alerts triggered by table and user statistics will not function until the **runstats** utility is run. This utility ensures that table and user statistics are up-to-date. When you define a scan job for DB2 databases, select the **Perform runstats on all tables being scanned** check box on the **Options** page.


**Note:** (DB2 only) When the following conditions apply, an alert is not generated:

- The triggering condition selected is Database-Tablespace Free Space Less Than.
- The storage entities selected are DB2 databases.
- The table spaces affected by the triggering condition are System Managed Space (SMS) tablespaces.

The alert is not triggered because IBM Tivoli Storage Productivity Center considers SMS containers to be directories on a file system

To create a database-table space alert, complete the following steps:

1. From the Navigation Tree pane, expand **Data Manager for Databases > Alerting**.
2. Right-click **Database-Tablespace Alerts** and click **Create Alert..** The Create Alert window opens and the Alert page is displayed.
3. Select the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** area, specify the event that will trigger the alert:
  - a. In the **Rdbms Type** field, select the type of database application. You can choose **All Rdbms**, **Oracle**, **SQL\Server**, **Sybase**, or **UDB**.
  - b. In the **Condition** field, select the triggering condition.
  - c. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be unavailable.
  - d. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be unavailable.
6. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see “Triggered actions for alerts” on page 1172.
7. Click the **Databases-Tablespaces** tab; the **Databases-Tablespaces** window opens.
8. Select the storage entities to which you want the alert to apply:
  - a. In the **Available** field, click the storage entities.


- b. Click . The storage entities are moved to the **Current Selections** field.
9. Save the alert:
  - a. Click **File > Save**.
  - b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

## Creating table alerts

You can use these alerts to monitor tables for certain conditions. When the condition is detected during a scan, an alert is generated. These alerts can be applied to table groups and tables only.

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services > Configuration > Alert Disposition**.
- Tables must be grouped into a table group before they can be included in an alert.


To create a table alert, complete the following steps:

1. From the Navigation Tree pane, expand **Data Manager for Databases > Alerting**.
2. Right-click **Table Alerts** and click **Create Alert..** The Create Alert window opens and the Alert page is displayed.
3. Select the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** area, specify the event that will trigger the alert:
  - a. In the **Rdbms Type** field, select the type of database application. You can choose **All Rdbms**, **Oracle**, **SQL\Server**, **Sybase**, or **UDB**.
  - b. In the **Condition** field, select the triggering condition.
  - c. In the **Value** field, type a numeric value. Depending on the triggering condition that you selected, this field might be unavailable.
  - d. In the **Value Units** field, select the unit of measurement. Depending on the triggering condition that you selected, this field might be unavailable.
6. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see “Triggered actions for alerts” on page 1172.
7. Click the **Tables** tab; the **Tables** window opens.
8. Select the storage entities to which you want the alert to apply:
  - a. In the **Available** field, click the storage entities.
  - b. Click . The storage entities are moved to the **Current Selections** field.
9. Save the alert:
  - a. Click **File > Save**.
  - b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

## Creating fabric alerts

You can use these alerts to monitor fabrics for changes. When the condition is detected during a probe, an alert is generated. These alerts can be applied to fabrics, zones, and zone sets.

If necessary, complete the following prerequisites before you begin the task:

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
  - If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services > Configuration > Alert Disposition**.
1. In the Navigation Tree pane, expand **Fabric Manager > Alerting**.
  2. Right-click **Fabric Alert** and click **Create Fabric Alerts**. The Create Fabric Alerts window opens and the Alert page is displayed.
  3. Select the **Enabled** check box.
  4. In the **Description** field, type a brief description of the job.
  5. In the **Triggering Condition** area, specify the event that will trigger the alert.
  6. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see “Triggered actions for alerts” on page 1172.
  7. Click the **Fabrics** tab; the Fabrics page opens.
  8. Select the storage entities to which you want the alert to apply:
    - a. In the **Available** field, click the storage entities.
    - b. Click . The storage entities are moved to the **Current Selections** field.
  9. Save the alert:
    - a. Click **File > Save**.
    - b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

## Creating storage system alerts

You can use storage systems alerts to monitor systems for changes in configuration or status. When the condition is detected during a probe, an alert is generated. These alerts can be applied to storage subsystem groups and storage subsystems.

If necessary, complete the following prerequisites before you begin the task:


- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services > Configuration > Alert Disposition**.

To create a storage subsystem alert, complete the following steps:

1. Expand the **Disk Manager > Alerting** node in the navigation tree.
2. Right-click **Storage Subsystem Alerts**, and click **Create Storage Subsystem Alert**. The Create Storage Subsystem Alert window is opened.
3. In the **Description** field, type a brief description of the job.
4. Select the **Enabled** check box.

5. In the **Triggering Condition** pane, specify the event that will trigger the alert:
  - a. In the **Condition** field, select the triggering condition. See “Triggering conditions for alerts” on page 1158 for a list of triggering conditions that are available for each alert type.
  - b. If you selected a performance threshold, specify the thresholds:
    - 1) In the **Critical Stress** field, type the upper boundary for the threshold. When this boundary is exceeded, a critical condition is recognized.
    - 2) In the **Warning Stress** field, type the upper boundary for the threshold. When this boundary is exceeded, a warning condition is recognized.
    - 3) In the **Warning Idle** field, type the lower boundary for the threshold. When the value falls below this boundary, a warning condition is recognized.
    - 4) In the **Critical Idle** field, type the lower boundary for the threshold. When the value falls below this boundary, a critical condition is recognized.

For more information about thresholds, see “Performance thresholds” on page 1176

6. In the **Alert Suppression** pane, choose the condition under which alerts are triggered, and then choose how to manage repeating alerts. For an explanation of the options in this pane, see “Alert suppression actions” on page 1171.
7. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see “Triggered actions for alerts” on page 1172.
8. Click the **Storage Subsystems** tab. The Storage Subsystems page is displayed.
9. Select the storage systems and storage system groups to which you want the alert to apply:
  - a. In the **Available** field, click the storage entities.
  - b. Click . The storage entities are moved to the **Current Selections** field.
10. Save the alert:
  - a. Click **File > Save**.
  - b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

## Creating switch alerts

You can use these alerts to monitor switches for changes in configuration or status. When the condition is detected during a probe, an alert is generated. These alerts can only be applied to switches.

If necessary, complete the following prerequisites before you begin the task:

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services > Configuration > Alert Disposition**.

To create a switch alert, complete the following steps:


1. In the Navigation Tree pane, expand **Fabric Manager > Alerting**.

2. Right-click **Switch Alerts** and click **Create Switch Alerts**. The Create Switch Alerts page opens and the Alert page is displayed.
3. Select the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** pane, specify the event that will trigger the alert:
  - a. In the **Condition** field, select the triggering condition. See “Triggering conditions for alerts” on page 1158 for a list of triggering conditions that are available for each alert type.
  - b. If you selected a performance threshold, specify the thresholds:
    - 1) In the **Critical Stress** field, type the upper boundary for the threshold. When this boundary is exceeded, a critical condition is recognized.
    - 2) In the **Warning Stress** field, type the upper boundary for the threshold. When this boundary is exceeded, a warning condition is recognized.
    - 3) In the **Warning Idle** field, type the lower boundary for the threshold. When the value falls below this boundary, a warning condition is recognized.
    - 4) In the **Critical Idle** field, type the lower boundary for the threshold. When the value falls below this boundary, a critical condition is recognized.

For more information about thresholds, see “Performance thresholds” on page 1176

6. In the **Alert Suppression** pane, choose the condition under which alerts are triggered, and then choose how to manage repeating alerts. For an explanation of the options in this pane, see “Alert suppression actions” on page 1171.
7. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see “Triggered actions for alerts” on page 1172.
8. Click the **Switches** tab. The Switches page is displayed.

**Note:** If you selected a performance threshold as the triggering condition, only those switches for which performance data can be collected are listed.

9. Select the storage entities to which you want the alert to apply:
  - a. In the **Available** field, click the storage entities.
  - b. Click . The storage entities are moved to the **Current Selections** field.

**Note:** You can select one or more switches from the **Switches** category. For device triggering conditions only, you can select one or more fabrics from the **Switches In Fabric** category. Selecting fabrics will include all switches in those fabrics. You can also select from both categories.

10. Save the alert:
  - a. Click **File > Save**.
  - b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.


## Creating endpoint device alerts

You can use these alerts to monitor endpoint devices. When the condition is detected during a probe, an alert is generated. These alerts can be applied to endpoint devices only.

If necessary, complete the following prerequisites before you begin the task:

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services > Configuration > Alert Disposition**.

To create an endpoint device alert, complete the following steps:

1. In the Navigation Tree pane, expand **Fabric Manager > Alerting**.
2. Right-click **Endpoint Device Alerts** and click **Create Endpoint Device Alerts**. The Create Endpoint Device Alerts window opens and the Alert page is displayed.
3. Select the **Enabled** check box.
4. In the **Description** field, type a brief description of the job.
5. In the **Triggering Condition** area, specify the event that will trigger the alert.
6. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see “Triggered actions for alerts” on page 1172.
7. Click the **Endpoint Devices** tab; the Endpoint Devices page opens.
8. Select the storage entities to which you want the alert to apply:
  - a. In the **Available** field, click the storage entities.
  - b. Click . The storage entities are moved to the **Current Selections** field.
9. Save the alert:
  - a. Click **File > Save**.
  - b. In the Save As window, type a name for the alert in the **Specify Alert name** field and click **OK**. The alert is displayed in the navigation tree.

### Database connection alerts

Database connection alerts notify the Tivoli Storage Productivity Center administrator when the database connection used by the Tivoli Storage Productivity Center server is lost unexpectedly and when the database connection is restored.

If necessary, complete the following prerequisites before you begin the task:

- If you want a script to be run after an event has triggered an alert, you must create the script before you create the alert.
- If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services > Configuration > Alert Disposition**.

The database connection alerts are always enabled and configured to add the alert to the alert log and send an email to the data administrator.

The database connection alerts are:

- DB connection failed
- DB connection successful

You cannot create, configure, delete, disable, edit, or rename these alerts. You can, however, change the destination for the database connection alerts. Expand **Administrative Services > Configuration**. Left-click **Alert Disposition**. In the right pane, under Email, you can change all of the fields.

## Configuring alert destinations

Entries in the Alert Disposition window define destination host addresses for SNMP, IBM Tivoli Enterprise Console, and email.

If you want to send SNMP traps, IBM Tivoli Enterprise Console events, or email when an alert is triggered, you must configure IBM Tivoli Storage Productivity Center with the alert destinations and email settings.

To configure alert dispositions, complete the following steps:

1. In the Navigation Tree pane, expand **Administrative Services > Configuration**.
2. Click **Alert Disposition**. The Edit Alert Disposition window opens.
3. In the **SNMP** area, specify the information about where SNMP traps are sent:
  - a. In the **Community** field, type the name of the SNMP community. By default, this is set to public.
  - b. In the **Host** field, type the fully qualified name of the system that you have configured to receive SNMP traps.
  - c. In the **Port** field, type the port number to which Tivoli Storage Productivity Center will direct SNMP traps. By default, this is set to 162.
  - d. If you want to configure a second SNMP destination, type the required information in the fields in the **SNMP Destination 2** area.
4. In the **TEC** area, specify information about where Tivoli Enterprise Console events are sent:
  - a. In the **Server** field, type the fully qualified name of the Tivoli Enterprise Console server.
  - b. In the **Port** field, type the port number to which Tivoli Storage Productivity Center will direct Tivoli Enterprise Console events. By default, this is set to 5529.
5. In the **Email** area, configure email settings:
  - a. In the **Mail Server** field, type the name of the mail server.
  - b. In the **Mail Port** field, type the port number for the SMTP server.
  - c. In the **Default Domain** field, type the domain, for example, MyCompany.com.
  - d. In the **Return to** field, type the email address to which undeliverable mail is sent.
  - e. In the **Reply to** field, type the email address to which email is sent when someone replies to an email notification about an alert. (It is a best practice to specify the email address of an Tivoli Storage Productivity Center administrator.)
  - f. In the **Data Administrator** field, type the email address of the Data Administrator.
6. In the **Alert Log Disposition** field, select the length of time to keep entries in the alert logs. The default is 90 days.



## Deleting alerts

To ensure that the list of alerts in the navigation tree is up-to-date, you can delete alerts that you no longer want to implement.

To delete an alert, complete the following steps:

1. In the navigation tree, complete one of the following actions:

| Alert type                     | Action                                                                                  |
|--------------------------------|-----------------------------------------------------------------------------------------|
| For computer alerts            | Expand <b>Data Manager &gt; Alerting &gt; Computer Alerts.</b>                          |
| For file system alerts         | Expand <b>Data Manager &gt; Alerting &gt; Filesystem Alerts.</b>                        |
| For directory alerts           | Expand <b>Data Manager &gt; Alerting &gt; Directory Alerts.</b>                         |
| For hypervisor alerts          | Expand <b>Data Manager &gt; Alerting &gt; Hypervisor Alerts</b>                         |
| For other NAS alerts           | Expand <b>Data Manager &gt; Alerting &gt; Other NAS Alerts</b>                          |
| For instance alerts            | Expand <b>Data Manager for Databases &gt; Alerting &gt; Instance Alerts.</b>            |
| For database-tablespace alerts | Expand <b>Data Manager for Databases &gt; Alerting &gt; Database-Tablespace Alerts.</b> |
| For table alerts               | Expand <b>Data Manager for Databases &gt; Alerting &gt; Table Alerts.</b>               |
| For storage subsystem alerts   | Expand <b>Disk Manager &gt; Alerting &gt; Storage Subsystem Alerts.</b>                 |
| For fabric alerts              | Expand <b>Fabric Manager &gt; Alerting &gt; Fabric Alerts.</b>                          |
| For switch alerts              | Expand <b>Fabric Manager &gt; Alerting &gt; Switch Alerts.</b>                          |
| For endpoint device alerts     | Expand <b>Fabric Manager &gt; Alerting &gt; Endpoint Device Alerts.</b>                 |

The list of existing alerts is displayed.

2. Right-click the alert and click **Delete**. The Delete Alert window opens.

**Tip:** If **Delete** is not available, you do not have authority to delete the alert.

3. Click **Yes**.

## Disabling an alert

You can disable an alert. This retains the alert definition but prevents the alert from being run.

To disable an alert, complete the following steps:

1. In the Navigation Tree pane, complete one of the following actions:

|                        |                                                                  |
|------------------------|------------------------------------------------------------------|
| For computer alerts    | Expand <b>Data Manager &gt; Alerting &gt; Computer Alerts.</b>   |
| For file system alerts | Expand <b>Data Manager &gt; Alerting &gt; Filesystem Alerts.</b> |
| For directory alerts   | Expand <b>Data Manager &gt; Alerting &gt; Directory Alerts.</b>  |
| For hypervisor alerts  | Expand <b>Data Manager &gt; Alerting &gt; Hypervisor Alerts.</b> |

|                                |                                                                                         |
|--------------------------------|-----------------------------------------------------------------------------------------|
| For other NAS alerts           | Expand <b>Data Manager &gt; Alerting &gt; Other NAS Alerts.</b>                         |
| For instance alerts            | Expand <b>Data Manager for Databases &gt; Alerting &gt; Instance Alerts.</b>            |
| For database-tablespace alerts | Expand <b>Data Manager for Databases &gt; Alerting &gt; Database-Tablespace Alerts.</b> |
| For table alerts               | Expand <b>Data Manager for Databases &gt; Alerting &gt; Table Alerts.</b>               |
| For storage subsystem alerts   | Expand <b>Disk Manager &gt; Alerting &gt; Storage Subsystem Alerts.</b>                 |
| For fabric alerts              | Expand <b>Fabric Manager &gt; Alerting &gt; Fabric Alerts.</b>                          |
| For switch alerts              | Expand <b>Fabric Manager &gt; Alerting &gt; Switch Alerts.</b>                          |
| For endpoint device alerts     | Expand <b>Fabric Manager &gt; Alerting &gt; Endpoint Device Alerts.</b>                 |

The list of existing alerts is displayed.

2. Click the alert that you want to disable. The Edit Alert page opens.
3. Clear the **Enabled** check box.
4. Click **File > Save.**

**Related topics:**

“Alerts” on page 109

## Editing alerts

You can edit an alert if you want to change the triggering condition, triggered action, or the storage resources against which it is deployed.

To edit an alert, complete the following steps:

1. In the navigation tree, complete one of the following actions:

| Alert type                     | Action                                                                                  |
|--------------------------------|-----------------------------------------------------------------------------------------|
| For computer alerts            | Expand <b>Data Manager &gt; Alerting &gt; Computer Alerts.</b>                          |
| For file system alerts         | Expand <b>Data Manager &gt; Alerting &gt; Filesystem Alerts.</b>                        |
| For directory alerts           | Expand <b>Data Manager &gt; Alerting &gt; Directory Alerts.</b>                         |
| For hypervisor alerts          | Expand <b>Data Manager &gt; Alerting &gt; Hypervisor Alerts</b>                         |
| For other NAS alerts           | Expand <b>Data Manager &gt; Alerting &gt; Other NAS Alerts</b>                          |
| For instance alerts            | Expand <b>Data Manager for Databases &gt; Alerting &gt; Instance Alerts.</b>            |
| For database-tablespace alerts | Expand <b>Data Manager for Databases &gt; Alerting &gt; Database-Tablespace Alerts.</b> |
| For table alerts               | Expand <b>Data Manager for Databases &gt; Alerting &gt; Table Alerts.</b>               |
| For storage subsystem alerts   | Expand <b>Disk Manager &gt; Alerting &gt; Storage Subsystem Alerts.</b>                 |
| For fabric alerts              | Expand <b>Fabric Manager &gt; Alerting &gt; Fabric Alerts.</b>                          |

|                            |                                                                         |
|----------------------------|-------------------------------------------------------------------------|
| For switch alerts          | Expand <b>Fabric Manager &gt; Alerting &gt; Switch Alerts.</b>          |
| For endpoint device alerts | Expand <b>Fabric Manager &gt; Alerting &gt; Endpoint Device Alerts.</b> |

The list of existing alerts is displayed.

2. Click the alert that you want to edit. The Edit Alert page opens.
3. Make any necessary changes.
4. Click **File > Save.**

**Tip:** Some alerts are shown automatically when a client application is registered to receive alerts from Tivoli Storage Productivity Center. You can identify these alerts by the keyword *Pre-defined* in their names. You can hide these alerts from view in the GUI. For more information, see the IBM Tivoli Storage Productivity Center Problem Determination Guide.

## Renaming alerts

You can rename alerts. This might be useful if the original name is ambiguous or lacks clarity.

To rename an alert, complete the following steps:

1. In the navigation tree, complete one of the following actions:

| Alert type                     | Action                                                                                  |
|--------------------------------|-----------------------------------------------------------------------------------------|
| For computer alerts            | Expand <b>Data Manager &gt; Alerting &gt; Computer Alerts.</b>                          |
| For file system alerts         | Expand <b>Data Manager &gt; Alerting &gt; Filesystem Alerts.</b>                        |
| For directory alerts           | Expand <b>Data Manager &gt; Alerting &gt; Directory Alerts.</b>                         |
| For hypervisor alerts          | Expand <b>Data Manager &gt; Alerting &gt; Hypervisor Alerts</b>                         |
| For other NAS alerts           | Expand <b>Data Manager &gt; Alerting &gt; Other NAS Alerts</b>                          |
| For instance alerts            | Expand <b>Data Manager for Databases &gt; Alerting &gt; Instance Alerts.</b>            |
| For database-tablespace alerts | Expand <b>Data Manager for Databases &gt; Alerting &gt; Database-Tablespace Alerts.</b> |
| For table alerts               | Expand <b>Data Manager for Databases &gt; Alerting &gt; Table Alerts.</b>               |
| For storage subsystem alerts   | Expand <b>Disk Manager &gt; Alerting &gt; Storage Subsystem Alerts.</b>                 |
| For fabric alerts              | Expand <b>Fabric Manager &gt; Alerting &gt; Fabric Alerts.</b>                          |
| For switch alerts              | Expand <b>Fabric Manager &gt; Alerting &gt; Switch Alerts.</b>                          |
| For endpoint device alerts     | Expand <b>Fabric Manager &gt; Alerting &gt; Endpoint Device Alerts.</b>                 |

The list of existing alerts is displayed.

2. Right-click the alert, and click **Rename.** The Rename Alert window opens.

3. Type the new name, and click **OK**.

## Viewing alerts

This section describes different methods for viewing alerts.

To view an alert, choose one of the following methods:

### Alert log

Click **IBM Tivoli Storage Productivity Center > Alerting > Alert Log**. Click **All Alerts** or any node that is highlighted in red and preceded by a red circle. Depending on what you select, the Alert History page displays all alerts or only the alerts that you selected. Figure 29 shows the alert history page.

|  | Computer                | Obj. Type              | Object                  | Alert Type                 | First Triggered         |
|--|-------------------------|------------------------|-------------------------|----------------------------|-------------------------|
|  | (N/A)                   | Configuration Analysis | db2admin.sqz5d3         | Policy Violation           | Apr 23, 2007 9:30:32 AM |
|  | (N/A)                   | Configuration Analysis | db2admin.sqz5d3         | Policy Violation           | Apr 23, 2007 9:30:32 AM |
|  | (N/A)                   | Discovery              | TPCUser.CIMOM Discovery | Job Failed                 | Apr 23, 2007 3:39:59 AM |
|  | (N/A)                   | Discovery              | TPCUser.CIMOM Discovery | Job Failed                 | Apr 22, 2007 3:39:58 AM |
|  | WORKGROUP\smartblued840 | Discovery              |                         | Entity Discovered          | Apr 22, 2007 2:29:56 AM |
|  | (N/A)                   | Discovery              | TPCUser.CIMOM Discovery | Job Failed                 | Apr 21, 2007 3:41:52 AM |
|  | bubs.sanjose.ibm.com    | (N/A)                  |                         | Performance monitor failed | Apr 20, 2007 1:31:06 PM |

Figure 29. Alert History page

To view details about each alert, click or double-click . The Detail for Alert page displays additional information (see Figure 30 on page 399). You can obtain similar information about any alert by clicking **Topology** to access the topology viewer. Details are available through the *alert overlay* in the graphical view and the Alert tab in the table view. The following figure shows the Detail for Alert page, which lists the name of the job, number of the run, type of alert, state of the alert, timestamp, name of the alert creator, name of the alert, and text about the alert

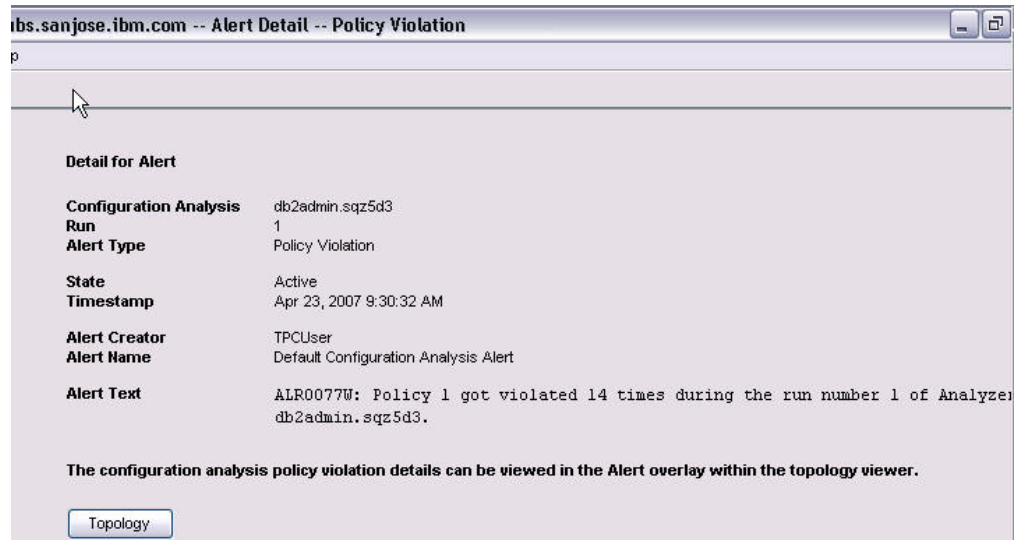


Figure 30. Detail for Alert page

### Alert overlay

The alert overlay represents various types of alerts. You can enable or disable the alert overlay in the graphical view of the topology viewer by right-clicking a blank space in the background of the graphical view, clicking **Global Settings** from the pop-up menu, and clicking **Alert** under Active overlays (see Figure 31). After you make your selection, click **OK**.

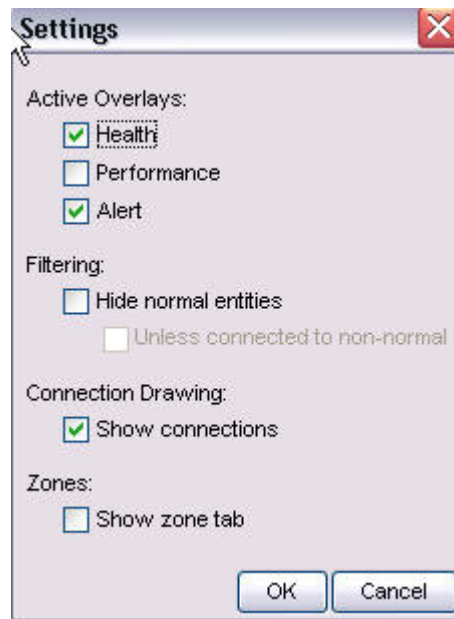



Figure 31. Enabling or disabling the alert overlay. In this example, the alert overlay is enabled.

When the alert overlay is active, entities that have alerts associated with them display a red alert overlay icon  that is located to the right of the entity and below the pinning icon. The icon is similarly located in the title bar of groups and tabs. Note that in the Configuration History page the alert overlay is disabled.

The alert overlay provides rollup, similar to pinning. If alerts are inside a collapsed group, the number of alerts within that group are aggregated and indicated in the collapsed group with an alert icon. All alerts are aggregated into one count, independent of type or entity involved. Hover over the alert icon for a tab or group to determine the total number of alerts that are rolled up into that icon and the time stamp of the newest alert.

In the topology viewer, right-click an entity with an alert icon to display the pop-up menu. Click **Show Alerts in Tabular View** to move to the Alert tab and to highlight alerts that are associated with that entity in the table view (see Figure 32). When you select an alert under the Alert tab, entities that are associated with that alert are highlighted in purple in the graphical view. Such an action might highlight more than one entity in the topology viewer.

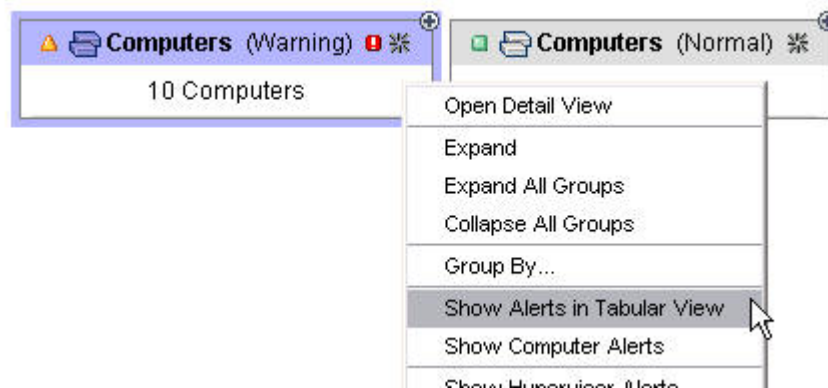


Figure 32. Using the pop-up menu to view alerts in the table view

External alerts do not display in the alert overlay. These alerts are usually Common Information Model (CIM) indications that cannot be tied to specific resources. To view the list of external alerts, click **IBM Tivoli Storage Productivity Center > Alerting > Alert Log > External**.

### Alert tab

The Alert tab is located in the table view of the topology viewer (see Figure 33 on page 401). Click this tab to view a hierarchical list of all current alerts that are associated with entities shown in the topology. For configuration analysis alerts, the top level provides information about which policy was violated during a particular analyzer run. The next level indicates the number of times that the policy was violated. This list contains one node for each policy violation. The Alert tab only lists the policy violations that contain at least one of the affected entities that are displayed in the graphical view. The remaining policy violations are filtered. The only exception is the Overview panel where the Alert tab lists all of the Tivoli Storage Productivity Center alerts. Click each violation to view a list of affected entities. Parent (enclosing) entities display in a separate column.

| Action: Expand All Groups |                          | Locate                  |                                                             |
|---------------------------|--------------------------|-------------------------|-------------------------------------------------------------|
| Computer                  |                          | Alert                   |                                                             |
| ...                       | Condition                | △ Last Triggered        | Enclosing Entity                                            |
| ⊖                         | <b>Policy 1 Violated</b> | Apr 23, 2007 9:30:32 AM |                                                             |
| ⊖                         | <b>Violation 1</b>       |                         |                                                             |
|                           | Violation 1              |                         | Fabric (200000000000B000)                                   |
|                           | Violation 1              |                         | Computer (hydrogen.almaden.ibm...FCPort (A00000000000A003)) |
| ⊕                         | <b>Violation 2</b>       |                         |                                                             |
| ⊕                         | <b>Violation 3</b>       |                         |                                                             |
| ⊕                         | <b>Policy 2 Violated</b> | Apr 23, 2007 9:30:32 AM |                                                             |

Figure 33. Viewing alerts by using the Alert tab in the table view. Expand the levels to view information about each violation.

For information about fixing policy violations, see “Resolving policy violations” on page 414. You can resort the alert list by clicking the headers in the table. To find entities associated with an alert in the graphical view, click the alert in the table to highlight (in purple) the corresponding entity and any associated entities.

## Viewing alert logs

The alert logs record information about the generated alerts. There are two alert logs that you can view:


- For alerts generated by Data Manager for Databases, expand **Data Manager for Databases > Alerting > Alert Log**. You can view:
  - All alerts or alerts organized by RDBMS such as by Oracle, SQL Server, Sybase, and UDB. If you do not license for a specific RDBMS, no nodes for that RDBMS are displayed under the **Alert Log** node.
  - Alerts by type (such as by Instance, Database, Tablespace, Table, User)
  - Alerts directed to you (the user currently logged on to the UI component)
- For alerts generated by all other components, expand **IBM Tivoli Storage Productivity Center > Alerting > Alert Log**. The alert types that are displayed depend on the Tivoli Storage Productivity Center license type. The following is a list of the possible alert types that are displayed under **Alert Log**:
  - All
  - Alerts Directed to <user-id> - Alerts directed to the user logged on to the interface.
  - Authentication Configuration
  - Replication
  - Storage Subsystem
  - Computer
  - Disk
  - Filesystem
  - Directory
  - User
  - OS User Group
  - Fabric
  - Switch
  - Endpoint Device
  - External
  - Tape Library
  - Configuration Analysis
  - Hypervisor

For alert types that are not displayed for certain license types, see “Licenses for Tivoli Storage Productivity Center” on page 14

Each alert node has an icon next to it that indicates its status:

- Red-outlined circle: a negative alert has occurred, such as job failed, new disk defect.
- Green-outlined square: a neutral alert has been detected, such as a new disk detected.

When you click an alert node, the alert log window opens a table containing

information for all the alerts for that node. Click  on a row to drill down for more information.

The following table describes the buttons at the top of the Alert Log window:

| Button     | Description                                                                                                                                                                                                    |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Clear      | Clear the selected alerts. When you click this button, the icon beside the alert changes to indicate that the alert has been cleared. The alert continue to be displayed in the alert log.                     |
| Delete     | Delete the selected alerts from the alert log.                                                                                                                                                                 |
| Clear ALL  | Clear all of the alerts in the report. When you click this button, the icon beside all the alerts changes to indicate that the alerts have been cleared. The alerts continue to be displayed in the alert log. |
| Delete ALL | Delete all of the alerts in the current report from the alert log.                                                                                                                                             |
| Refresh    | Regenerate the report you are viewing to pick up any newly generated alerts that have been triggered since you last generated the report.                                                                      |

**Tip:** Some alerts are shown automatically when a client application is registered to receive alerts from Tivoli Storage Productivity Center. You can identify these alerts by the keyword *Pre-defined* in their names. You can hide these alerts from view in the GUI. For more information, see the IBM Tivoli Storage Productivity Center Problem Determination Guide.

## Viewing tape library alerts

View discovered alert conditions in the tape libraries. Alerts are based on the CIM\_AlertIndication specification of SMI-S.

To view these alerts, expand **IBM Tivoli Storage Productivity Center > Alerting > Alert Log** and click the **Tape Library** node.

---

## Checking storage configurations

This section introduces methods that a system administrator can use to assess changes that have occurred in SAN configurations over time and to check whether SAN configurations comply with predefined policies.

## Using configuration history

This topic describes how system administrators can capture and analyze historical data that identifies possible problems with a storage area network (SAN) configuration.



**Note:** The configuration history feature is available only with the IBM Tivoli Storage Productivity Center Standard Edition.

The configuration history feature takes and displays snapshots of changes that occurred in your SAN configuration over a period of time that you specify. After you set the time period (how often to take snapshots and how long to store them), in a page similar to the topology viewer you can manipulate a snapshot selection panel to show changes that occurred between two or more points in the time period. System administrators can use the configuration history feature to:

#### **Correlate performance statistics with configuration changes**

For example, during collection of performance statistics (including volume performance statistics) on an TotalStorage Enterprise Storage Server system you might delete a volume. While no new statistics are reported on that volume, the Tivoli Storage Productivity Center Performance Manager would have already collected partial statistical information prior to the deletion. At the end of the data collection task, reporting of the partially collected statistics on the (now) deleted volume would require access to its properties which would not be available. The configuration history feature, with its ability to take and store snapshots of a system's configuration, could provide access to the volume's properties.

#### **Analyze end-to-end performance**

You want to know why performance changed on volume A during the last 24 hours. To learn why, it is useful to know what changes were made to the storage subsystem's configuration that might affect the volume's performance, even if performance statistics were not recorded on some of those elements. For example, even if performance statistics on a per-rank basis are not collected, but the number of volumes allocated on a rank is increased from 1 to 100 over time, access to that configuration history information helps with analyzing the volume's degraded performance over time.

#### **Aid in planning and provisioning**

The availability of configuration history can enhance the quality of both provisioning and planning. For example, historical data is useful when using the Tivoli Storage Productivity Center Volume Performance Advisor to provision a volume or when using the Tivoli Storage Productivity Center Version 3.3 Integrated Host, Security and Subsystem Planner to plan tasks.

To use the configuration history feature, complete the following steps:

1. In the Navigation tree pane, expand **Administrative Services > Configuration > Configuration History Settings**. The Configuration History Settings page displays for you to indicate how often to capture SAN configuration data and how long to retain it.
2. Perform the following to collect historical data:
  - a. In the **Create snapshot every** field, click the check box to enable this option and type how often (in hours) you want the system to take snapshot views of the configuration.
  - b. In the **Delete snapshots older than** field, click the check box to enable this option and type how long (in days) you want the snapshots to be stored.
  - c. The page displays the total number of snapshots in the database and the date and time of when the latest snapshot was taken. To refresh this information, click **Update**.

- d. To optionally create and title a snapshot on demand, in the **Title this snapshot** field type a name for the on demand snapshot and click **Create Snapshot Now**. If you do not want to title the on demand snapshot, simply click **Create Snapshot Now**.
  - e. To return to the default settings (the default settings are create snapshots every 12 hours and delete snapshots older than 14 days), click **Reset to defaults**.
  - f. To save your settings, click **File > Save**.
3. In the Navigation tree pane, expand **Tivoli Storage Productivity Center > Analytics > Configuration History**, and click **Configuration History**. The software loads the snapshot data for the length of time that you specified. The Configuration History page (a variation of the topology viewer) displays the configuration's entities and a floating snapshot selection panel. The panel allows you to define the time periods against which the configuration is compared to determine whether changes have occurred.

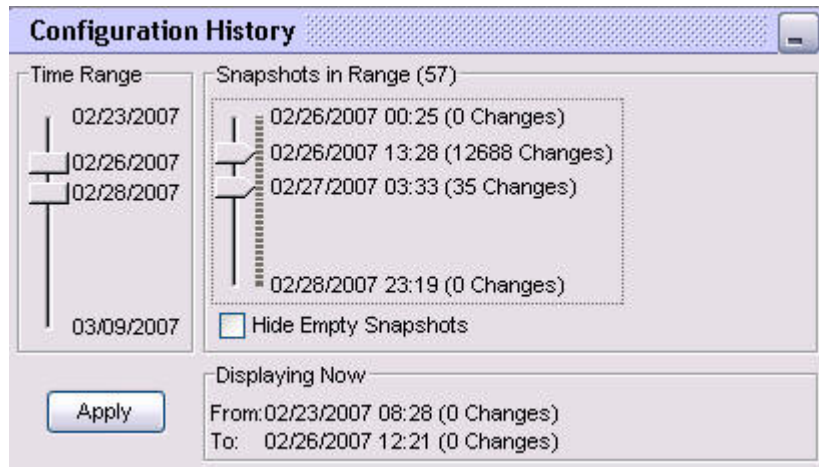


Figure 34. Snapshot selection panel. Use the thumb sliders to establish the time interval that you want to examine.

4. To define the time periods that you want to compare, perform the following:
  - a. Using the mouse, drag the two "thumbs" in the left **Time Range** slider to establish the desired time interval. The Time Range slider covers the range of time from the oldest snapshot in the system to the current time. It indicates the date as *mm/dd/yyyy*, where *mm* equals the month, *dd* equals the day, and *yyyy* equals the year.
  - b. Drag the two thumbs in the right **Snapshots in Range** slider to indicate the two snapshots to compare. The Snapshots in Range slider allows you to select any two snapshots from the time interval specified by the Time Range slider. The value in parentheses beside the Snapshots in Range slider indicates the total snapshots in the currently selected time range. The Snapshots in Range slider has one tick mark for each snapshot from the time interval that you specified in the Time Range slider. Each snapshot in the Snapshots in Range slider is represented as time stamp *mm/dd/yyyy hh:mm*, where the first *mm* equals the month, *dd* equals the day, *yyyy* equals the year, *hh* equals the hour, and the second *mm* equals the minute. The value in parentheses beside each snapshot indicates the number of changes that have occurred between this and the previous snapshot. Snapshots with zero changes are referred to as empty snapshots. If you provided a title while creating an on demand snapshot, the title displays after the time

stamp. If you want to remove empty snapshots, click the check box to display a check mark in **Hide Empty Snapshots**. The **Displaying Now** box indicates the two snapshots that are currently active.

- c. Click **Apply**.
5. Determine the changes that have occurred to the entities by examining the icons and colors associated with them in the graphical and table views: For information about viewing the changes, see “Viewing configuration changes in the graphical view” and “Viewing configuration changes in the table view” on page 407.

One single snapshot selection panel applies for all Configuration History views that are open at the same time. Any change that you make in this panel is applied to all of the Configuration History views.

**Important:** If you manually deleted a missing entity from the Tivoli Storage Productivity Center, the configuration history feature rediscovers and displays the entity in the configuration history viewer as both Missing and in a Normal state. For example, Entity A (with ID 12) is missing and you decide to manually delete it by using the Remove From Database option from the pop-up menu for Entity A in the topology viewer or configuration history viewer. Later, the Tivoli Storage Productivity Center rediscovers Entity A and assigns it a new ID (ID 22). The configuration history feature treats the same Entity A (with ID 12 and ID 22) as two different entities and displays them in both Missing and Normal states in the configuration history viewer. This scenario can also occur if a missing entity is removed by the Tivoli Storage Productivity Center removed resource retention function and is rediscovered by the Tivoli Storage Productivity Center again at a later time.

### Viewing configuration changes in the graphical view

This section describes how to determine the changes that have occurred to a configuration over time by examining the icons and colors of the change overlay in the graphical view of the Configuration History page.

In the Configuration History page, a *change overlay* presents icons and colors that indicate changes in the configuration between the time that a snapshot was taken and the time that a later snapshot was taken:

- The icons display beside the name of the entity in the graphical view and in the change status column of the table view.
- The colors display as background colors for entities in the graphical view and for affected rows in the table view.

Table 19 describes the icons and colors of the change overlay.

Table 19. Icons and colors of the change overlay





| Change overlay icon                                                                                      | Change overlay color indicator | Description                                                                                                      |
|----------------------------------------------------------------------------------------------------------|--------------------------------|------------------------------------------------------------------------------------------------------------------|
| <br>Yellow pencil     | Blue background                | Entity changed between the time that the snapshot was taken and the time that a later snapshot was taken.        |
| <br>Light gray circle | Dark gray background           | Entity did not change between the time that the snapshot was taken and the time that a later snapshot was taken. |

Table 19. Icons and colors of the change overlay (continued)

| Change overlay icon                                                                                 | Change overlay color indicator | Description                                                                                                              |
|-----------------------------------------------------------------------------------------------------|--------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| <br>Green cross    | Green background               | Entity was created or added between the time that the snapshot was taken and the time that a later snapshot was taken.   |
| <br>Red minus sign | Red background                 | Entity was deleted or removed between the time that the snapshot was taken and the time that a later snapshot was taken. |
| Not applicable                                                                                      | Light gray background          | Entity did not exist at the time that the snapshot was taken or at the time that a later snapshot was taken.             |

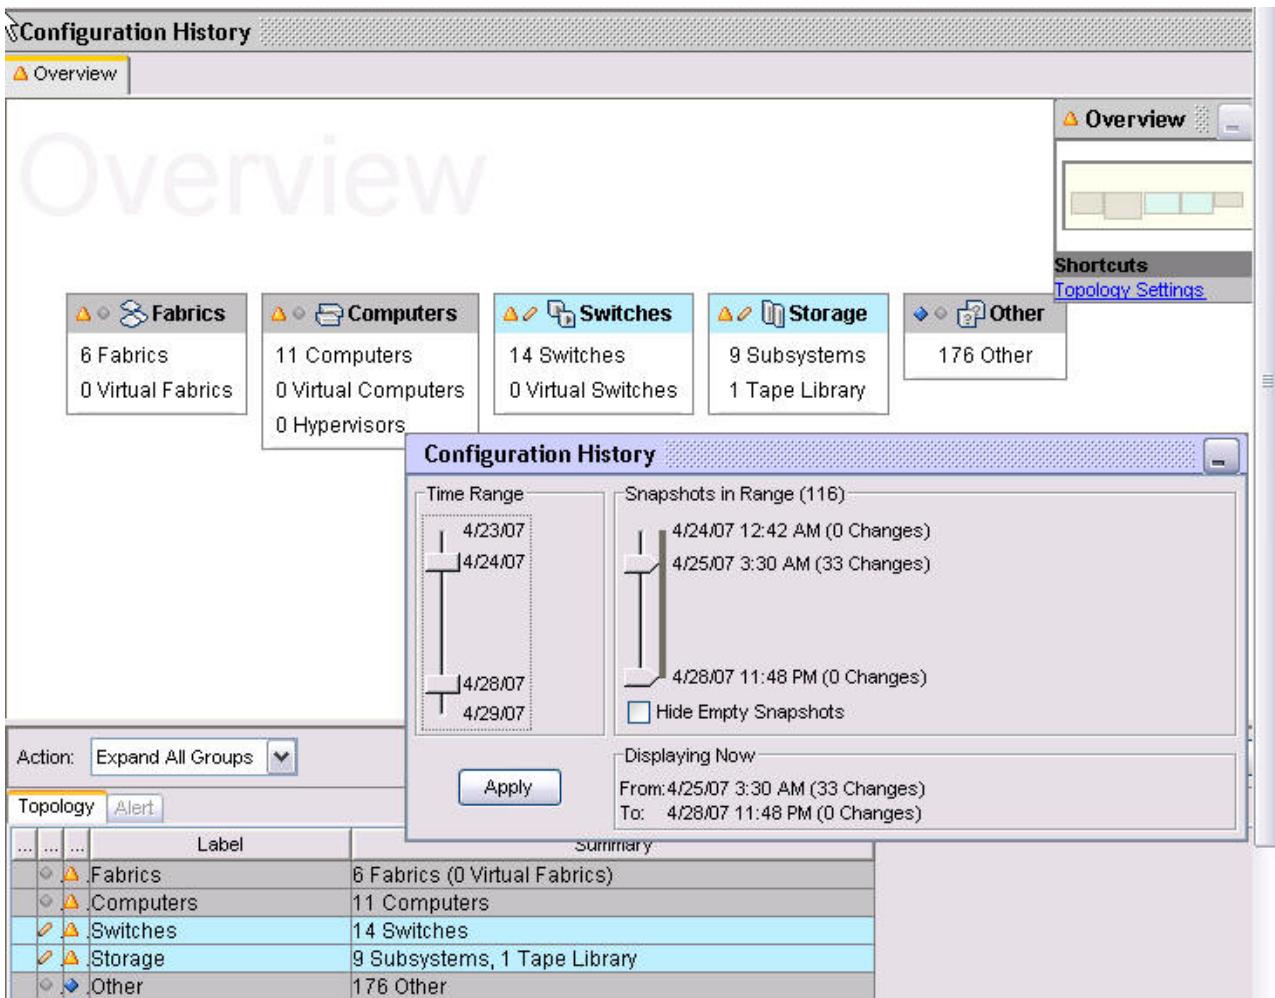


Figure 35. Icons and colors of the change overlay. In the graphical view, the pencil icon beside the switches and storage entities and the blue background color indicate that change occurred to these entities. The pencil icon and blue background also appears for these entities in the table view. In the snapshot selection panel, use the Time Range and Snapshots in Range sliders to determine when the change occurred.

To distinguish them from tabs in the topology viewer page, tabs in the Configuration History page (Overview, Computers, Fabrics, Storage, and Other) have a light gray background and are outlined in orange. The minimap in the Configuration History page uses the following colors to indicate the aggregated change status of groups:

- **Blue** - One or more entities in the group have changed. Note that the addition or removal of an entity is considered a change.
- **Gray** - All of the entities in the group are unchanged.

Entities in the graphical view may be *active* (they existed at one or both snapshots) or *inactive* (not yet created or deleted). *Active entities* act like they normally do in the topology viewer; when you select them all relevant information also appears in the table view. You can adjust a grouping of active entities, but you may not perform actions that change the database, such as pinning. *Inactive entities* do not exist in the selected snapshots, but exist in the other snapshots. They are shown in a light gray background and do not have a change icon associated with them. Inactive entities display to keep the topology layout stable and to make it easier to follow what has changed (instead of having the entities flicker in and out of existence when you change the snapshot selection). Inactive entities are not listed in the table view.

An entity that is moved from one group to another group appears only once in the new group in the graphical view. For example, if the health status of a computer has changed from Normal to Warning, the Configuration History page displays the computer as changed in the Warning health group (and no longer displays the computer in the Normal health group).

**Note:** In the Configuration History view, the performance and alert overlays are disabled and the minimap's shortcut to the Data Path Explorer is not available.

### Viewing configuration changes in the table view

This section describes how to determine the changes that have occurred to a configuration over time by examining the icons and colors of the change overlay in the table view of the Configuration History page.

In the table view, a changed cell or tab is color coded. The change overlay uses the same colors and icons for the table view as for the graphical view (see Table 19 on page 405). The changed cell or tab displays both the old and new values as **Old Value > New Value**. To display a summary of change for the entire row, hover over the changed row. The summary displays the title of the column. It also displays the old and new value for each cell or tab that changed. For example, if a group contains one new entity but everything else in the group is unchanged, the group as a whole displays as "changed." If you click the group, the one new entity displays as "created" and all other entities display as "unchanged."

The Change column in the table view also indicates changed entities.


An object that is moved from one group to another group appears only once in the new group in the graphical view. For example, if the health status of a computer has changed from Normal to Warning, the Configuration History page displays the computer as changed in the Warning health group (and no longer displays the computer in the Normal health group).

## Using configuration analysis

This topic describes how system administrators can determine whether an existing SAN configuration complies with predefined best practices (policies).

**Note:** The configuration analysis feature is available only with the IBM Tivoli Storage Productivity Center Standard Edition.

The configuration analysis feature allows a system administrator to select up to 16 policies and to specify the name of a target SAN (or zone set) for the analysis. The analysis tool reads the specified SAN data and the policy information for the specified policies from the IBM Tivoli Storage Productivity Center database. It outputs the policy violations as Tivoli Storage Productivity Center alerts. To view an alert, you can use one of three methods:

- Click the alert in the Configuration Analysis alert log
- Click the alert overlay icon (a red exclamation mark ) next to the affected entity
- Click the Alert tab in the table view of the topology viewer

System administrators can use configuration analysis to learn of any best-practices violations their actions may have caused. For example, in provisioning a new storage subsystem, an administrator may inadvertently violate a zoning-related policy. If the zoning-related configuration changes persist in the database, configuration analysis can detect this policy violation.

To perform configuration analysis, complete the following steps:

1. Ensure that you have previously run discovery and probe jobs for the computers, fabrics, switches, storage, and other objects of interest.
2. Ensure that the Tivoli Storage Productivity Center Data Server and Device server are running. To do so, in the navigation tree pane, expand the following:
  - **Administrative Services > Services > Data Server > Server**
  - **Administrative Services > Services > Device Server > Device server**
3. Ensure that the following Tivoli Storage Productivity Center agents are installed and configured:

**Important:** Policy violations are not raised if the agent information needed to evaluate a policy is not collected and stored in the Tivoli Storage Productivity Center database.

- A mix of in-band or Storage Resource agent, out-of-band agent, and CIMOM fabric agents to collect all fabric topology information (switches, interconnections, and zoning). This information is needed for the evaluation of all configuration analysis policies.
- CIMOM agents for storage subsystems to collect the storage subsystem information (subsystem details, storage volumes, storage ports and storage volumes to hosts' mappings). This information is needed for the evaluation of configuration analysis policies 2, 3, 5, 7 and 11. CIMOMs are not required to communicate with storage systems that use native interfaces. Tivoli Storage Productivity Center uses native interfaces to communicate with IBM System Storage DS8000 Storage Manager, IBM XIV Storage System, and IBM System Storage SAN Volume Controller.
- Data agents or Storage Resource agents to collect the host information (operating system type and version). This information is needed for the evaluation of configuration analysis policies 3 and 9.
- In-band Fabric agents or Storage Resource agents to collect the attribute information from the host information (HBA and host port details). This information is needed for the evaluation of configuration analysis policies 2, 3, 4, 8 and 9.
- CIMOM agents for Tape to collect the tape information. This is needed for the evaluation of configuration analysis policy 2.

4. In the navigation tree pane, expand **IBM Tivoli Storage Productivity Center > Analytics**.
5. Right-click **Configuration Analysis**. A selection list displays.
6. Click **Create Analyzer**. The Create Analyzer window opens (see Figure 36).

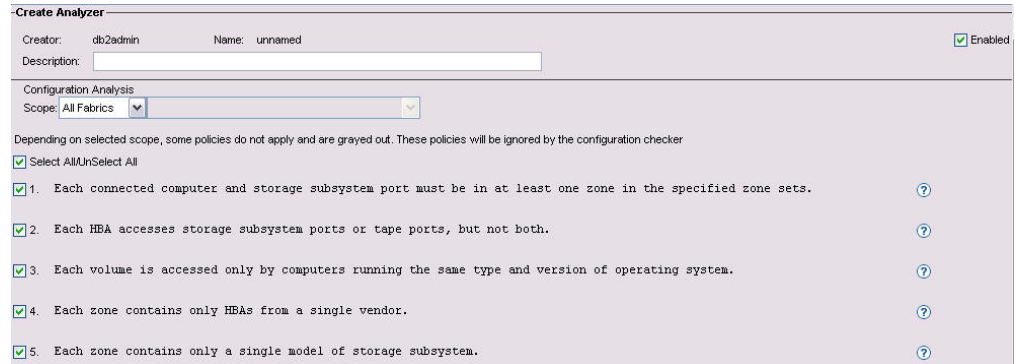


Figure 36. The Create Analyzer window. The figure shows a partial view of the window. Use the Create Analyzer window to select the scope and the policies of the configuration analysis.

7. Perform the following to define the analysis:
  - a. Click the check box to display a check mark beside **Enabled**.
  - b. In the **Description** field, type a brief description of the of the analysis job.
  - c. To specify the scope of the SAN data to be checked, make a selection from the **Configuration Analysis Scope** list. Choices include:
    - All Fabrics
    - One fabric
    - One Zoneset
 If you select One fabric or One Zoneset, an adjacent list displays for you to click a specific fabric or zone set.
  - d. Check the SAN data against up to 16 policies by performing one of the following:
    - Choose all of the policies by clicking the check box to display a check mark beside **Select All/Unselect All**.
    - Choose one or more individual policies by clicking the check box to display a check mark beside each policy.

**Note:** For Policies 12 and 13, type values up to a maximum of 9999.


Table 20 lists the policies and their explanations. You can also access each explanation by clicking  at the right of the page.

Table 20. Policies used by the configuration analysis feature

| Policy                                                                                                         | Explanation                                                                                                                                                                                                                                                                                                     |
|----------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1. Each connected computer and storage subsystem port must be in at least one zone in the specified zone sets. | Determines whether an administrator forgot to zone a connected port. Putting connected ports into zones is useful for security and performance reasons. Ports are usually grouped into zones based on applications, server operating systems, or HBA vendors. The Fabric scope is not supported by this policy. |

Table 20. Policies used by the configuration analysis feature (continued)

| Policy                                                                                                      | Explanation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2. Each HBA accesses storage subsystem ports or tape ports, but not both.                                   | Determines whether an HBA accesses both storage subsystem and tape ports. Because HBA buffer management is configured differently for storage subsystems and tape, it is not desirable to use the same HBA for both disk and tape traffic. The Fabric and zone set scopes are not supported by this policy because an HBA can be connected to multiple fabrics.                                                                                                                                                                     |
| 3. Each volume is accessed only by computers running the same type and version of operating system.         | Determines whether computers that run different operating systems access the same storage volumes. Use of the same volumes by computers that run different operating systems may corrupt the data that is stored on the volumes. This applies, regardless of whether the computers are in the same zone.                                                                                                                                                                                                                            |
| 4. Each zone contains only HBAs from a single vendor.                                                       | Determines whether HBAs from different vendor types are in their own zone. Receiving a registered state change notification (RSCN) may cause an HBA to lose a zoned device, preventing the HBA from seeing or communicating with other devices in the zone. To avoid losing a zoned device, keep HBAs from different vendor types in their own zone.                                                                                                                                                                                |
| 5. Each zone contains only a single model of storage subsystem.                                             | Determines whether different storage subsystems are in the same zone. While no technical problem is associated with storage subsystems from different vendors and of different models being in the same zone, an administrator may find them more difficult to organize. When similar storage systems are in the same zone, an administrator can easily group them for different applications.                                                                                                                                      |
| 6. Each zone is part of a zone set.                                                                         | Determines the presence of orphan zones. Orphan zones are not associated with any zone set. They are not useful because their definitions are not used and they take up switch resources.                                                                                                                                                                                                                                                                                                                                           |
| 7. Each host must be zoned so that it can access all of its assigned volumes.                               | Determines whether the zones that were configured by the storage administrator allow each computer to access all of the storage volumes that are assigned to it. The administrator specifies the storage subsystem ports through which the computer port accesses volumes, but may forget to configure zones that enable the ports to communicate during volume assignment. This policy also determines whether zoning makes assigned volumes inaccessible to the computer ports. The Fabric scope is not supported by this policy. |
| 8. Each computer has only HBAs of the same model and firmware version.                                      | Checks whether there is only one type of HBA in each computer. Using only one type of HBA minimizes configuration problems. The policy also checks whether firmware upgrades have been done properly for all HBAs in a computer. The zone set scope is not supported by this policy.                                                                                                                                                                                                                                                |
| 9. For each host type and operating system, every HBA of a given model must have the same firmware version. | Determines whether all firmware upgrades have been done for the HBAs in the operating system. The zone set scope is not supported by this policy.                                                                                                                                                                                                                                                                                                                                                                                   |



Table 20. Policies used by the configuration analysis feature (continued)

| Policy                                                                                                                                                           | Explanation                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 10. Every SAN switch of a given model must have the same firmware version.                                                                                       | Determines whether firmware upgrades have been done for all switches of the same type. For example, if you have four identical models of SAN switches from the same vendor and you perform a firmware upgrade on one, it is best to perform the upgrade on all of the others. The zone set scope is not supported by this policy.                                                                                                                                |
| 11. Every storage subsystem of a given model must have the same firmware version.                                                                                | Determines whether firmware upgrades have been done for all storage subsystems of the same type. For example, if you have four identical storage subsystems from the same vendor and you perform a firmware upgrade on one, it is best to perform the upgrade on all of the others. The zone set scope is not supported by this policy.                                                                                                                          |
| 12. Each fabric may have a maximum of x zones.                                                                                                                   | Checks whether the number of zone definitions in the fabric is larger than the number that you entered. In large fabrics, too large a number of zone definitions can become a problem. Fabric zone definitions are controlled by one of the switches in that fabric, and limiting their number ensures that the switch's zoning tables do not run out of space. The zone set scope is not supported by this policy. You can enter up to a maximum of 9999 zones. |
| 13. Each zone may have a maximum of x zone members.                                                                                                              | Checks whether the number of zone members in a zone is larger than the number that you entered. In large fabrics, too large a number of zone members can become a problem. Fabric zone members are controlled by one of the switches in that fabric, and limiting their number ensures that the switch's zoning tables do not run out of space. You can enter up to a maximum of 9999 zones.                                                                     |
| 14. Replication Plan is intact with respect to the SRG(s) and the replication session associated during planning through SAN Planner.                            | Determines whether SRGs that have associated resiliency profiles are intact with respect to their members and their replication relationships.                                                                                                                                                                                                                                                                                                                   |
| 15. All the source volumes involved in Metro Mirror Failover/Failback sessions are conforming to 1:4 primary to secondary LSS for a failback direction scenario. | Metro Mirror (Failover/Failback) can be directional, and up to 4 LSS secondaries can be connected to a primary LSS.                                                                                                                                                                                                                                                                                                                                              |
| 16. Inter/intra site connectivity is valid for replication plan deployments.                                                                                     | Checks whether connectivity either direct or through fiber channel fabric exists for resiliency session deployments.                                                                                                                                                                                                                                                                                                                                             |

8. To specify how often to run the analysis, scroll to **Scheduling**.
9. In the **How often to run** area, schedule how often you want the job to run. Choose from the following actions:

**Run Now**

Run the job immediately.

**Run Once at**

Run the job once at a specified date and time. Use the list boxes to specify the month, day, and year. Type the hour and minutes over the existing defaults. Click AM or PM.

**Note:** The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval

between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

### Run Repeatedly

Run the job repeatedly according to a schedule that you specify:

- Indicate when to begin the job. Use the list boxes to specify the month, day, and year. Type the hour and minutes over the existing defaults. Click AM or PM.
  - Indicate how often to run the job. Select one of the following:
    - In the **Repeat Every** field, click the desired number of minutes, hours, days, weeks, or months.
    - In the **Run on these days** field, click to display a check mark next to the desired days.
10. In the **How to handle time zones** area, indicate the time zone to be used when scheduling and running the policy. Click one of the following:

#### Use the time zone that the server runs in

Select this option to use the time zone of the location where the Data server resides.


#### Use this time zone

Select this option to indicate a specific time zone for when to run the job. From the list box, select the time zone.

11. From the menu bar, select **File > Save**. The Save As pop-up opens.
12. Type a name for the analyzer and click **OK**. The Notice pop-up displays **Analyzer submitted**.
13. Click **OK** to begin the job. In the Navigation tree pane, the newly named analyzer job displays under **IBM Tivoli Storage Productivity Center > Analytics > Configuration Analysis**. The format of the name is *creator.analyzer*, where *creator* is the name that displays near the top of the Create Analyzer window and *analyzer* is name that you assigned in step 12 (for example *db2admin.analyzer2*).
14. Expand the job to view information about the run. The information includes the number of times that the job has been run, and the date and time of the run. The color of the type indicates the status:
- Blue indicates that the run is in progress.
  - Green indicates that the run is finished.
- To refresh the status, right-click the job and click **Update Job Status**.
15. Click the run to view information about it. On a page titled with the name of the job and run number, the following information displays:

| Field              | Description                                                                       |
|--------------------|-----------------------------------------------------------------------------------|
| <b>Run</b>         | Number of the run                                                                 |
| <b>Status</b>      | The status of the run (Success or Failed), not whether policy violations occurred |
| <b>Start Time</b>  | The time at which the run started                                                 |
| <b>Finish Time</b> | The time at which the run ended                                                   |
| <b># Jobs</b>      | The number of jobs                                                                |
| <b># failed</b>    | The number of jobs that failed                                                    |
| <b>Computer</b>    | The IP address of the computer on which the job was run                           |



|                      |                                        |
|----------------------|----------------------------------------|
| <b>Status</b>        | The status of the run                  |
| <b>Start Time</b>    | The time at which the run started      |
| <b>Finish Time</b>   | The time at which the run ended        |
| <b>Log File Name</b> | The complete path name of the log file |

16. To view whether a policy violation occurred, click . The Job log file page displays the following information:



- Scope of the run (all fabrics, one fabric, or one zone set)
- Policies that were checked
- Total violations
- Breakdown of violations per policy
- Whether the run completed successfully or the errors it encountered



**Note:** To avoid overloading the Topology Viewer, the analysis job stores and displays only the first 50 policy violations. You must resolve these violations and run the job again to view any remaining violations, or you can change this policy violation limit by using the `setdscfg` command. The attribute to set is `MaxPolicyViolationsStored`.

17. To view one or more alerts generated by policy violations, in the Navigation tree expand **IBM Tivoli Storage Productivity Center > Alerting > Alert Log > Configuration Analysis**. The Alert History - All Policy Violations page displays a log of job runs that generated alerts for policy violations. A policy violation alert is generated for each policy that was violated during a run. A policy may be violated several times, but only one alert is generated; the text in the alert indicates the number of times the policy was violated. The Alert History - All Policy Violations page includes the following information:

| Column                                                                              | Description                                                                                                     |
|-------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
|  | When clicked, provides details about the alert.                                                                 |
|  | When clicked, makes the Clear and Delete buttons active. When double-clicked, provides details about the alert. |
| <b>Computer</b>                                                                     | Not applicable (N/A)                                                                                            |
| <b>Configuration Analysis</b>                                                       | The name of the configuration analysis job                                                                      |
| <b>Alert Type</b>                                                                   | The type of alert generated during the run                                                                      |
| <b>First Triggered</b>                                                              | The date and time of when the first alert was generated                                                         |
| <b>Last Triggered</b>                                                               | The date and time of when the last alert was generated                                                          |
| <b># Times</b>                                                                      | The number of times that an alert was generated for a run                                                       |
| <b>Alert Creator</b>                                                                | The userid of the person who created the run                                                                    |
| <b>Alert Name</b>                                                                   | The official name of the alert                                                                                  |
| <b>Script</b>                                                                       | Not applicable (N/A)                                                                                            |

For more information about the alert log, including the buttons at the top of the page, see “Viewing alert logs” on page 401.

18. To view details about each alert, in the Alert History - All Policy Violations page click the run that you want, then click  or double-click . The Detail for Alert page displays the following information:

| Field                  | Description                                                                                                                                                                                                                                                                                                                                          |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configuration Analysis | The name and scope of the analyzer job.                                                                                                                                                                                                                                                                                                              |
| Run                    | The number of run in which the alert occurred.                                                                                                                                                                                                                                                                                                       |
| Alert Type             | The type of alert (Policy Violation).                                                                                                                                                                                                                                                                                                                |
| State                  | The status of the alert. An <i>active</i> (unresolved) state displays a red exclamation mark<br><br>; an <i>acknowledged</i> (resolved) state displays a green check mark<br><br>. |
| Timestamp              | The date and time that the alert was created. Expressed as the month, day, year, hour, minute, second, and AM or PM.                                                                                                                                                                                                                                 |
| Alert Creator          | The userid of the person who created the analyzer job.                                                                                                                                                                                                                                                                                               |
| Alert Name             | The name of the alert.                                                                                                                                                                                                                                                                                                                               |
| Alert Text             | A description of the policy that was violated and the number of times that it was violated during that analyzer run.                                                                                                                                                                                                                                 |

You can also use the topology viewer to view alerts about policy violations. To launch the viewer, click **Topology**. For more information, see “Viewing alerts” on page 398.

### Resolving policy violations

When you click the Alert tab in the topology viewer and expand the entries in the table view, information displays about each configuration analysis policy violation. To verify and fix the policy violation, use the information in the following table.

| Configuration Analysis Policy                                                                                         | Enclosing Entity              | Affected Entity | How to Verify and Fix the Policy Violation                                                                                                                                                                                                                                          |
|-----------------------------------------------------------------------------------------------------------------------|-------------------------------|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Policy 1. Each connected computer and storage subsystem port must be in at least one zone in the specified zone sets. | Computer, Subsystem, or Other | Port            | In the zone tab within the L2:Fabric view, verify that the port is not a member of the zone set. Fix the violation by adding the port to the appropriate zone within the zone set.                                                                                                  |
|                                                                                                                       | Fabric                        | Zone set        |                                                                                                                                                                                                                                                                                     |
| Policy 2. Each HBA accesses storage subsystem ports or tape ports, but not both.                                      | Computer                      | HBA             | In the zone tab within the L2:Computer view, verify that HBA ports are accessing both the disk and tape subsystems. Fix the violation by configuring the computer access to disk and tape in two separate zone sets, and ensuring that only one of these is active at a given time. |
| Policy 3. Each volume is accessed only by computers running the same type and version of operating system.            | Subsystem                     | Volume          | In L1:Computers view, one can see that the computers accessing the volume do not have the same OS type and version. This violation can be fixed by mapping or assigning the volume only to the computers with same OS type and version.                                             |
|                                                                                                                       | Computer                      | Port            |                                                                                                                                                                                                                                                                                     |
|                                                                                                                       | Computer                      | Port            |                                                                                                                                                                                                                                                                                     |

| Configuration Analysis Policy                                                                                      | Enclosing Entity  | Affected Entity | How to Verify and Fix the Policy Violation                                                                                                                                                                                                                                                                                         |
|--------------------------------------------------------------------------------------------------------------------|-------------------|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Policy 4. Each zone contains only HBAs from a single vendor.                                                       | Fabric            | Zone            | In the zone tab within the L2:Fabric view, find the computers in the zone. Open the respective L2:Computer views and select the HBA tab in the table view to view the vendor of the HBAs on those computers. Fix the violation by editing the zone definition and ensuring that all the HBAs in the zone are from the same vendor. |
| Policy 5. Each zone contains only a single model of storage subsystem.                                             | Fabric            | Zone            | In the L2:Fabric view, verify that the models of the subsystems in the zone are different. Fix the violation by editing the zone definition and ensuring that all the subsystems in the zone have the same model.                                                                                                                  |
|                                                                                                                    | Subsystem         | Port            |                                                                                                                                                                                                                                                                                                                                    |
| Policy 6. Each zone is part of a zone set.                                                                         | Fabric            | Zone            | In the zone tab within the L2:Fabric view, verify that the zone does not belong to any of the zone sets and is listed under Orphan Zones. Fix the violation by adding the zone to a zone set or by deleting the zone.                                                                                                              |
| Policy 7. Each host must be zoned so that it can access all of its assigned volumes.                               | Computer or Other | Port            | None of the zones in the L2:Computer or L2:Subsystem views has both the computer and subsystem port. Fix the violation by adding or editing the zone definition, and by ensuring that both the computer and subsystem port are in the same zone.                                                                                   |
|                                                                                                                    | Subsystem         | Port            |                                                                                                                                                                                                                                                                                                                                    |
|                                                                                                                    | Subsystem         | Volume          |                                                                                                                                                                                                                                                                                                                                    |
| Policy 8. Each computer has only HBAs of the same model and firmware version.                                      | Computer          | Computer        | Open the L2:Computer view and select the HBA tab in the table view to verify that the model and firmware version for the HBAs on that computer are different. Fix the violation by installing the HBAs of the same model or upgrading the firmware version of the HBA on that computer.                                            |
| Policy 9. For each host type and operating system, every HBA of a given model must have the same firmware version. | Computer          | HBA             | Open the respective L2:Computer views and select the HBA tab in the table view to verify that the firmware and driver version for the HBAs on those computers are different. Fix the violation by upgrading the firmware and driver versions of the HBAs on those computers.                                                       |
|                                                                                                                    | Computer          | HBA             |                                                                                                                                                                                                                                                                                                                                    |
| Policy 10. Every SAN switch of a given model must have the same firmware version.                                  | Switch            | Switch          | Open the L1:Switches view and select the switch tab to verify that the firmware versions on those switches are different. Fix the violation by upgrading the firmware versions on those switches.                                                                                                                                  |
|                                                                                                                    | Switch            | Switch          |                                                                                                                                                                                                                                                                                                                                    |

| Configuration Analysis Policy                                                                                                                                           | Enclosing Entity             | Affected Entity         | How to Verify and Fix the Policy Violation                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Policy 11. Every storage subsystem of a given model must have the same firmware version.                                                                                | Subsystem                    | Subsystem               | Open the L1:Subsystems view and select the subsystem tab to verify that the firmware versions on those subsystems are different. Fix the violation by upgrading the firmware versions on those subsystems.                                                                                                                                                                                                                               |
|                                                                                                                                                                         | Subsystem                    | Subsystem               |                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Policy 12. Each fabric may have a maximum of $x$ zones.                                                                                                                 | Fabric                       | Fabric                  | In the zone tab within the L2:Fabric view, verify that the number of zones present in the fabric is greater than $x$ (the value entered by the user in this policy). Fix the violation by reducing the number of zones to a value less than $x$ .                                                                                                                                                                                        |
| Policy 13. Each zone may have a maximum of $x$ zone members.                                                                                                            | Fabric                       | Zone                    | Search for the zone in the zone tab in the L2:Fabric view. Verify that the number of zone members present in that zone is greater than $x$ (the value entered by the user in this policy). Fix the violation by reducing the number of zone members to a value less than $x$ .                                                                                                                                                           |
| Policy 14. Replication Plan is intact with respect to the SRG(s) and the replication session associated during planning through SAN Planner.                            | Storage Resource Group (SRG) | Replication session     | In the alert tab in the L0:Storage Resource Groups view, verify that the numbers of source and target volumes in the SRG are not equal or that the volumes are not in a replication relationship. Fix the violation by running the SAN Planner again and provisioning additional volumes or provisioning resiliency for the existing volumes.                                                                                            |
|                                                                                                                                                                         | SRG                          | Volume                  |                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Policy 15. All the source volumes involved in Metro Mirror Failover/Failback sessions are conforming to 1:4 primary to secondary LSS for a failback direction scenario. | SRG                          | Logical Subsystem (LSS) | IBM Tivoli Storage Productivity Center does not have an LSS view, but in the L2:Storage view, you can verify the details of the LSSs responsible for the violation. This violation only occurs if the copy session is created incorrectly in IBM Tivoli Storage Productivity Center for Replication. Fix the violation by changing the session details in Tivoli Storage Productivity Center for Replication to conform with the policy. |
| Policy 16. Inter/intra site connectivity is valid for replication plan deployments.                                                                                     | SRG                          | Replication session     | In the alert tab in the L0:Storage Resource Groups view, verify that there is no connectivity between source and target volumes. Fix the violation by manually providing a common fibre channel fabric or direct connection.                                                                                                                                                                                                             |

---

## Planning and modifying storage configurations

This section describes SAN Planner and how to use it to plan and modify changes to the system configuration.

### SAN Planner overview

SAN Planner assists the user in end-to-end planning involving fabrics, hosts, storage controllers, storage pools, volumes, paths, ports, zones, zone sets, storage resource groups (SRGs), and replication. Once a plan is made, the user can select to have the plan implemented by SAN Planner.

### Supported Systems

You can use the following products with SAN Planner:

- TotalStorage Enterprise Storage Server
- System Storage DS8000
- DS6000
- SAN Volume Controller
- Storwize V7000
- The XIV system

SAN Planner supports the **Space Only** workload profile option for any other storage system.

### Requirements for SAN Planner

For planning, IBM Tivoli Storage Productivity Center must be managing the host system, subsystem, and the fabric interconnecting them. If the host, subsystem, or fabric information is not collected by performing subsystem probes and fabric discovery, SAN Planner is not able to generate the plan or execute portions of the plan and issues an error message.

For volume or virtual disk creation, Tivoli Storage Productivity Center must be managing the subsystem. You must have run the subsystem probe either through an active CIM object manager (CIMOM) or native API connection as applicable for the subsystem.

For volume or virtual disk assignment, Tivoli Storage Productivity Center must be managing the host, fabric, and subsystem. The host and subsystem must be in the same fabric. If zoning changes are required, the SAN planner can be used to configure the zoning.

For fabric zoning configuration or fabric zone control operations, Tivoli Storage Productivity Center must manage the fabric. Depending on fabric type and agent capabilities, Tivoli Storage Productivity Center uses a CIM agent, inband agent, or Storage Resource Agent (SRA).

For a plan that requires a virtual disk to be assigned to a host, the virtual disk must have access to two nodes within an I/O group on the SAN Volume Controller or Storwize V7000. These two nodes must be shared on the same fabric as the host.

A storage controller performance monitor must be run to select the **Workload Profile** options. For SAN Volume Controller or Storwize V7000, it is required that

the performance data is collected for all its backend subsystems as well. If a performance monitor has not been run, the **Space Only** workload profile option is allowed for planning based on available capacity only.

The IBM Subsystem Device Driver (SDD) installed on a host is required for multipath planning.

**Note:**

- The Linux SDD driver is no longer available. Starting with Red Hat Enterprise Linux Version 5 and SUSE Linux Enterprise Server Version 10, only the DM\_Multipath is available. The DM\_Multipath is supported by SAN Planner
- The SDD driver is no longer supported on HP, starting with HP-UX 11i v3 with Itanium. HP-UX 11i v3 provides a native multipathing feature which is supported by SAN Planner.
- If you are using the SAN planning tools which collect data from Oracle Solaris systems, the HBAs must be configured to use Solaris systems. This note applies to Storage Resource agents.

Run the following command on the Solaris system:

```
cfgadm -al
```

If the system displays this message:

```
fc-fabric connected unconfigured unknown
```

run this command:

```
cfgadm -c configure c#
```

where *c#* is your HBA controller ID. An example of an HBA controller ID is: c5.

You must reboot the Solaris system.

- The SAN Planner does not support Easy Tier<sup>®</sup> (mixed pools) as input to a plan for automated storage provisioning. These pools are not displayed when you select resources for planning or request a list of recommended pools.
- When creating a plan with multipath provisioning, do not include computers that have EMC Powerpath installed. For example, do not include EMC Clariion volumes in a plan with multipath provisioning.

For replication planning, you must add all the subsystems to be used in planning to IBM Tivoli Storage Productivity Center for Replication. You must also ensure that connectivity is established between the subsystems.

## Creating a SAN plan

You can create a SAN plan by right clicking the **IBM Tivoli Storage Productivity Centre > Analytics > SAN Planner** node and selecting **Create Plan**. The planner wizard is started.

You can also create a SAN plan from the **IBM Tivoli Storage Productivity Centre > Disk Manager > SAN Planner** node and selecting **Create Plan**. The planner wizard is started.

The type of plan that you select determines the content of the remaining pages in the wizard. The options are:



**Provision storage only, including optional multipath and zone considerations**

Click to create a plan for provisioning new storage. The plan does not include provisioning for storage replication.

**Provision storage with replication, including optional multipath and zone considerations on the source resources**

Click to create a plan for provisioning new storage and storage replication.

**Provision replication to existing storage**

Click to create a plan for provisioning replication for existing storage (provision target volumes for existing source volumes).

**Configure multipathing or zoning only**

Click to create a plan for provisioning multipathing and zoning for existing storage.

To print plan information from the interface, click **File > Print** on the Tivoli Storage



Productivity Center menu bar, or click on the toolbar.

**Provision storage only, including optional multipath and zone considerations**

Use this option to provision storage volumes or virtual disks with multipathing and zoning. The input resources for planning can be storage subsystems, volumes, hosts, or storage resource groups (SRGs). If you are provisioning storage for SAN Volume Controller or Storwize V7000, you must select **Virtual Disks**.

An SRG is a set of entities managed by Tivoli Storage Productivity Center. These entities can be servers, switches, storage subsystems, fabrics, storage pools, and storage volumes. Using SRGs provides you with a means of separating your application data from your pool of unused resources. You can specify policies for provisioning (volume creation and selection, workload profiles, zoning, and multipathing configuration) and associate them with SRGs. These policies are used by SAN Planner to populate default settings.

You can use this option to plan and select appropriate storage controllers, storage pools, and storage volumes (when using unassigned volumes) that satisfy your inputs. This option enables you to select controller type preference, whether storage requests can be satisfied by multiple controller types, and the Redundant Array of Independent Disks (RAID) level. For SAN Volume Controller or Storwize V7000, the RAID level represents the RAID level for the backend controllers.

When creating new volumes on an extent pool, SAN Planner uses one of the following methods:

- If an encryption group is selected, SAN Planner creates volumes on pools that belong to the selected encryption group. Only pools that belong to that encryption group are used for planning. This option might not be enabled if the storage system does not support encryption.
- If an encryption group is not selected, SAN Planner uses the current performance utilization of storage resources to determine whether a new volume should be allocated on particular pool in a particular storage controller. If multiple storage pools from different controllers can potentially satisfy your provisioning request, then the SAN Planner uses the rated utilization of the

pools (the sum of the previous provisioning performance requirements, which might be greater than the current utilization) to break the ties and select a candidate storage pool.

You can also select other options in the SAN Planner such as selecting a workload profile as an input to the SAN Planner's analysis, excluding solid-state drives from planning, or using thin provisioning for creating thin provisioned volumes based on the thin provisioning profile selected.

See the **Configure multipathing or zoning only** section for details on SAN Planner's multipathing and zoning options.

## **Provision storage with replication, including optional multipath and zone considerations on the source resources**

In addition to provisioning storage with multipathing and zoning options, when you select replication options in the wizard, SAN Planner provisions replication for your storage resources. Resilient resource planning is available only for devices which are supported by IBM Tivoli Storage Productivity Center for Replication.

Depending on the storage resources you are planning for, you can select from the following copy services session types:

- FlashCopy
- VDisk mirroring
- Synchronous Metro Mirror failover/failback
- Synchronous Metro Mirror failover/failback with practice
- Synchronous Metro Mirror single direction
- Asynchronous Global Mirror either direction with two-site practice
- Asynchronous Global Mirror failover/failback
- Asynchronous Global Mirror failover/failback with practice
- Asynchronous Global Mirror single direction
- Three-site Metro Global Mirror with practice
- Three-site Metro Global Mirror
- Snapshot

For example, if you are planning for the XIV system, you can only select from Snapshot, Metro Mirror, or Global Mirror session types. Practice sessions are not supported for the XIV system.

## **Provision replication to existing storage**

This option allows you to extend existing replication sessions by adding more storage resources and protecting them. The new volumes are added to existing replication sessions. SAN Planner ensures that source volumes are added to the source side of a copy relationship and target volumes are added to the target side. For Snapshot copy sessions, SAN Planner does not provision or manage target volumes. The XIV system selects and manages the targets for Snapshot replication.

## **Configure multipathing or zoning only**

SAN Planner enables you to set up multipath options. The supported multipath driver on a host is the IBM SDD.

**Note:**

- The Linux SDD driver is no longer available. Starting with Red Hat Enterprise Linux Version 5 and SUSE Linux Enterprise Server Version 10, only the DM\_Multipath is available.
- The SDD driver is no longer supported on HP, starting with HP-UX 11i v3 with Itanium. HP-UX 11i v3 provides a native multipathing feature which is supported by SAN Planner.

SAN Planner enables system administrators and storage area network (SAN) administrators to plan and implement storage provisioning for hosts and storage systems with multipath support in fabrics managed by Tivoli Storage Productivity Center. Planning the provisioning of storage to hosts with multipath drivers requires knowing which storage systems are supported by the host multipath driver and the multipath modes that are supported by the driver and storage subsystem. Planning the paths between the hosts and storage controller requires designing paths between hosts and storage subsystems which are implemented through zones in the fabric.

SAN Planner is used for specifying multiple paths options between selected hosts and storage systems. There should be at least two paths available to a storage system and at least four paths available to SAN Volume Controller. Although not required, it is beneficial to have at least two host ports.

Path planning in SAN Planner does not directly interact with zone planning. The path planning option provides the path information which the zone planning option uses. Each path is represented by a host port WWPN, target port WWPN, and a Volume ID for the volume on the target that is mapped to the host port. These paths are created when the zone planning option is implemented.

SAN Planner specifies automatic zoning between ports on the selected hosts and subsystems in a fabric. All zoning is done with WWPN-based zoning. SAN Planner plans zoning configuration for new storage that is provisioned for use by a host. For example, the zoning option can be used when a new storage volume is created and assigned to a host. The zoning option can also be used with a volume that has already been created and is assigned to a host that requires more storage. In these cases, SAN Planner determines which host and storage need to be zoned together, which provides the zoning option with the exact set of ports that need to be zoned together. SAN Planner then uses the zoning inputs for the planning.

If the subsystem/host are within the same fabric and the zoning option is not checked, then existing zones or zone sets are used. If the zoning option is checked, a new zone or zone set is created.

To perform zoning, the host and subsystem being zoned by SAN Planner must be within the same fabric. For the case where the host and subsystem reside in more than one of the same fabrics, SAN Planner creates identical zones in all the fabrics.

The guidance policies used for zone planning are as follows:

- One zone per host bus adapter (HBA)
- One zone per host
- One zone per host port
- Auto zone: largest zone that satisfies the validation policies

The validation policies used for zone planning are as follows:

- Maximum number of zone members in a zone = N.
- No two controllers of different types should be in the same zone.
- Maximum number of zones in a fabric = N.

The second policy in the preceding list has a higher priority than the third policy.

## Planner and user inputs interdependencies

The planner options depend on the inputs that you select. For example, selecting a host that does not have an agent installed limits the planning that can be done. The planning process reveals inconsistencies in your preferences. For example, it might not be possible to have your specified number of paths and have a different zone for each host. In such cases, the planner makes a best effort to satisfy your inputs. If the planner cannot satisfy conditions, it issues warnings. Errors discovered before plan execution are displayed in the planner. You can adjust the inputs to correct the problem. Errors discovered during plan execution are in the job execution log.

## Using the SAN Planner

Learn how to use SAN Planner to select configuration information for a plan and modify a configuration.

You must run a fabric, disk and data discovery, and system and fabric configuration probe on storage controllers for planning purposes. You must also run a performance monitor on the storage controllers. You can use the following products with SAN Planner:

- TotalStorage Enterprise Storage Server
- System Storage DS6000
- System Storage DS8000
- SAN Volume Controller
- Storwize V7000
- The XIV system

SAN Planner supports the **Space Only** workload profile for any other storage system.

A Data agent or Storage Resource agent is required to obtain host information used in volume assignments. An in-band Fabric agent or Storage Resource agent is optional if the fabric information is available through another in-band Fabric agent or Storage Resource agent in the same fabric, an out-of-band agent, or through a CIM object manager (CIMOM) managed switch. Run the in-band Fabric agent or Storage Resource agent to get the host operating system information. If the in-band Fabric agent or Storage Resource agent was not run and the storage controller does not have information about the host operating system, an asterisk is displayed next to the host port. Ensure that the operating system-imposed limits on the number of volumes that can be accessed are not exceeded.

1. Expand the **IBM Tivoli Storage Productivity Center > Analytics or Disk Manager** node in the navigation tree.
2. Right-click **SAN Planner** and select **Create Plan**. The SAN Planner wizard is displayed. The wizard provides the interface for user input to put together a plan for volume, path, zone, and replication planning. It also provides the ability to get configuration recommendations based on the current storage configurations and the user inputs.

**Note:** Clicking a saved plan (located under the **SAN Planner** node in the navigation tree) opens the wizard with the Edit Plan *TPCuserid.plan\_name* pane instead of the Create Plan pane, where *TPCuserid* is the Tivoli Storage Productivity Center user ID and *plan\_name* is the name the plan was given when it was previously saved.

3. To move to the next page in the wizard, click **Next**.

### **SAN Planner wizard: Selecting the planning task**

Select the type of storage provisioning plan that you want to create. The type of plan that you select determines the content of the remaining pages in the SAN Planner Wizard.

1. Select the type of provisioning plan from the following options:

**Provision storage only, including optional multipath and zone considerations**

Click to create a plan for provisioning new storage. The plan does not include provisioning for storage replication.

**Provision storage with replication, including optional multipath and zone considerations on the source resources**

Click to create a plan for provisioning new storage and storage replication.

**Provision replication to existing storage**

Click to create a plan for provisioning replication for existing storage (provision target volumes for existing source volumes).

**Configure multipathing or zoning only**

Click to create a plan for provisioning multipathing and zoning for existing storage.

2. To move to the next page in the wizard, click **Next**.

### **SAN Planner wizard: Selecting the provision plan details**

Select the type of storage that you want to provision: volumes or virtual disks. You can also select whether you want to include multipath and zoning provisioning in the SAN plan.

1. To specify the type of storage that you want to provision, click **Volumes** or **Virtual Disks**. If you are provisioning storage for IBM System Storage SAN Volume Controller, you must select **Virtual Disks**.
2. Click **Include multipath planning in the plan** to enable multipath provisioning later in the wizard. This option is unavailable if you selected **Provision replication to existing storage** on the previous page of the wizard.
3. Click **Include fabric rezoning for host planning in the plan** to enable zoning provisioning later in the wizard. This option is unavailable if you selected **Provision replication to existing storage** on the previous page of the wizard.
4. To move to the next page in the wizard, click **Next**.


### **SAN Planner wizard: Specifying the plan content**

Select the input resources that you want to use in your Storage Area Network (SAN) plan. Input resources can be storage systems, volumes, hosts, or a storage resource group (SRG).

If you plan to use an SRG, you can select the SRG when you add the input resources at the top of this page, or you can select a candidate SRG from the **Use a candidate SRG** at the bottom of this page.

A candidate SRG is an SRG grouping storage resources together. By providing an SRG as a candidate SRG to SAN Planner, you are instructing SAN Planner to consider storage resources from this SRG when doing the capacity planning. Using a candidate SRG provides you with a means of separating your application data from your pool of unused resources.

You can create SRGs using the Storage Resource Group Management panel.

1. To specify the type of storage that you want to provision, click **Add**. The Topology Viewer in the Planner Selection window is opened. The Topology Viewer provides a topology view of the current system configuration.
2. Double-click the box title for a storage item in the topology viewer to expand the storage view into the L0 level view of the available systems. Before selecting systems, you might want to view the relationship between the systems and switch connectivity. To view relationships, double-click a specific system at the L0 Storage tab view. The L1 Subsystems view is opened, which displays the Subsystems and the Connectivity (Switches) boxes. Double-clicking a specific system displays the L2 view of that system and everything the system is connected to.
3. In the L0, L1 or L2 view, first expand the system using the + or – buttons in the upper right corner of the icon for the system. Then you must click the system itself to activate the >> button to include the system in the plan. Click each SAN Planner supported element in the same manner, that you want to include in the SAN plan. To select multiple elements, press and hold the **Ctrl** key and click each element icon.
4. Click  to move the selected elements into the **Selected Elements** pane.

**Tips:**

- The **Selected Elements** pane lists the storage resource groups, fabrics, computers, storage systems, pools, and volumes that are selected. When using the **Selected Elements** pane, the subsystem selections are not required to be in the same fabric. The **Fabrics** section lists the fabrics (by WWN) and the corresponding selected systems and selected hosts within each fabric. If a system has been configured in different fabrics, it is listed multiple times. The **Subsystems** section contains the selected storage systems. The number of fabrics to which the system belongs is shown next to each system. Selected pools are listed under their storage systems. Selected volumes are listed under their pools. The **Computers** section lists the selected hosts and the number of fabrics the host belongs to. The hosts are used in path planning and zone planning if **Include multipath planning in the plan** and **Include fabric rezoning for host planning in the plan** are selected in the Select Planning Tasks page.
  - To remove an element from the **Selected Elements** pane, select the element and click **Remove**.
5. Click **OK** to save your selections and return to the Specify Plan Content page where the storage elements that you selected are displayed. If an SRG is selected as an input resource, the name of the associated provisioning profile, if any, is displayed.
  6. To use a candidate SRG as a pool of storage resources from which to provision new storage, click **Use a candidate SRG**. Select the SRG that you want to use from the **Candidate SRG** list. If you select a candidate SRG, existing unassigned volumes belonging to the candidate SRG are used or new volumes are created from a candidate SRG. The name of the candidate SRG is displayed.

Unassigned volumes in the SRGs selected as input resources in the **Selected Elements** pane are not used and new volumes are not created from the input SRG.

7. If you have selected a candidate SRG, click **View SRG** to view the SRG details.
8. To move to the next page in the wizard, click **Next**.

**Related tasks:**

“Working with storage resource groups” on page 503

Use storage resource groups to organize logically related storage entities into named groups. For example, a storage resource group that represents a business critical email application might include all the hosts that participate in the email application cluster, the storage subsystems that provide storage to the application, and the switches and fabrics through which the application’s data is configured to travel.

### **SAN Planner wizard: Specifying selected volume assignments for input resources**

Select how the SAN Planner uses any individual unassigned volumes contained in the plan input resources.

If one or more of the volumes in your plan input resources is assigned to hosts, but not all of them are, then you can decide how to assign them.

1. Select the type of volume assignment for your input resources from the following options:

**Use these volumes as candidates for capacity planning**

Click to use the unassigned volumes as candidates for provisioning storage on the Specify Capacity Plan page.

**Assign these volumes to the hosts and skip capacity planning**

Click to omit the Specify Capacity Plan page and continue on with the next steps in the planner. You can view the volume assignment on the Review User Selections page.

**Ignore these volumes in the capacity planning**

Click to remove the volumes from consideration in the plan. If there are no storage pools or storage systems in the input resources, then no capacity planning can be done.

2. To move to the next page in the wizard, click **Next**.

### **SAN Planner wizard: Specifying selected volume assignments for an SRG**

Select how the SAN Planner uses any unassigned volumes contained in a candidate storage resource group (SRG).

If one or more of the volumes in your plan candidate SRG is assigned to hosts, but not all of them are, then you can decide how to assign them.

1. Select the type of volume assignment for your candidate SRG from the following options:

**Assign these volumes to the hosts**

Click to assign the volumes to all the hosts and continue to the Specify Capacity Plan page. You can view the volume assignment on the Review User Selections page.

**Ignore these volumes in the capacity planning**

Click to remove the volumes from consideration in the plan and continue to the Specify Capacity Plan page.

- To move to the next page in the wizard, click **Next**.

### SAN Planner wizard: Specifying capacity plan

Create and configure new storage volumes or virtual disks specified in your SAN plan.

- If you want to select any predefined workload profile other than **Space only**, you must have run a performance monitor on the storage controller.
- If you want to select a user-defined workload profile, you must first create a workload profile using the Workload Profiles panel.

Some of the field names that are displayed on the Specify Capacity Plan page differ depending on the type of storage (volumes or virtual disks) that you selected on the Select Provision Plan Details page.

- In the **Total Capacity** field, enter the total capacity of the volumes or virtual disks that are to be used for provisioning storage in GB. The Topology Viewer in the Planner Selection window is opened. The Topology Viewer provides a topology view of the current system configuration.
- In the **Number of Volumes** or **Number of Virtual Disks** field, enter the number of volumes or virtual disks that you want to create.
- In the **RAID level** field, select an Array of Independent Disks (RAID) level for the volume or virtual disk.
- In the **Volume Name Prefix** or **Virtual Disk Name Prefix** field, enter the prefix to be used in the volume or virtual disk name. If the prefix is longer than 14 characters, a warning prompt is displayed. On the plan recommendation page, if the SAN Planner determines that the prefix length exceeds the naming limitation of the storage system, you are prompted to change the prefix. The following values represent the maximum volume name length for IBM storage systems. The maximum volume name length includes the prefix that you enter + the volume name.

Table 21. Maximum volume name lengths for storage systems

| Storage system                                                                                                                                                                                                                                                                       | Maximum volume name length (in characters)* |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------|
| ESS                                                                                                                                                                                                                                                                                  | 8                                           |
| DS4000, DS5000                                                                                                                                                                                                                                                                       | 30                                          |
| DS6000, DS8000                                                                                                                                                                                                                                                                       | 16                                          |
| XIV                                                                                                                                                                                                                                                                                  | 63                                          |
| SAN Volume Controller 5.1 and earlier                                                                                                                                                                                                                                                | 15                                          |
| SAN Volume Controller 6.1                                                                                                                                                                                                                                                            | 60                                          |
| <b>Note:</b> Tivoli Storage Productivity Center adds one or more characters to the volume name depending upon how many volumes you are creating. For example, if 1-9 volumes are created, one character is added; if 10-99 volumes are created, two characters are added, and so on. |                                             |

- In the **Workload Profile** field, select a workload profile for the volume or virtual disk. You can use all predefined profiles other than **Space only** for the following products:
  - TotalStorage Enterprise Storage Server
  - System Storage DS6000
  - System Storage DS8000
  - SAN Volume Controller



- Storwize V7000
- the XIV system

With SAN Planner, you can use the **Space only** workload profile option for other storage systems.

6. For **The recommended provisioning will be based on the performance data collected between** section, enter a **Start Date** and **End Date** for the collection of performance data that you want to use in planning.
7. Click **Use unassigned volumes not involved in copy relationship** if you want to use existing unassigned volumes. Unassigned volumes that are in a copy relationship are excluded. Exception: The following unassigned volumes are excluded regardless of their copy relationship:
  - Unassigned volumes on IBM storage systems that have copy relationships that were not created in IBM Tivoli Storage Productivity Center or IBM Tivoli Storage Productivity Center for Replication.
8. To move to the next page in the wizard, click **Next**.

### **SAN Planner wizard: Specifying advanced capacity plan**

Select additional options for creating and configuring new storage volumes or virtual disks specified in your SAN plan.

1. To create thin volumes or virtual disks for IBM System Storage SAN Volume Controller, IBM System Storage DS8000, or IBM XIV Storage System, select one of the predefined profiles from the **Thin Provisioning Profile** field. You can also create a profile in the Thin Provisioning Profiles panel. The predefined options are as follows:

#### **Ignore thin provisioning**

A thin provisioning profile is not used.

#### **Standard**

This profile sets the over allocation limit for a pool at 200%. This percentage sets the virtual capacity at up to two times the real capacity.

#### **Conservative**

This profile sets the over allocation limit for a pool at 150%. This percentage sets the virtual capacity at up to one and a half times the real capacity.

#### **Aggressive**

This profile sets the over allocation limit for a pool at 300%. This percentage sets the virtual capacity at up to three times the real capacity.

2. To specify whether you want to use or exclude solid-state disks from provisioning, click one of the following options from the **Solid State Disks** field:

#### **Use Solid State Disks**

Click to use only solid-state disks for provisioning.

#### **Do not use Solid State Disks**

Click to exclude solid-state disks from provisioning,

3. To create new volumes on extent pools that belong to an encryption group, select a value from the **Encryption Group** list. If there are no encryption groups available, **<none selected>** is displayed in this box.
4. Click **Show Recommended Pools** to show the pools to be used for provisioning. The recommended pools are displayed under the corresponding storage systems in the **The following resources are selected for planning** box.

5. Optional: To remove pools from the provisioning process, select a pool or pools from the **The following resources are selected for planning** box and click **Remove Pool**.
6. To move to the next page in the wizard, click **Next**.

**Related tasks:**

“Working with thin provisioning profiles” on page 563

Learn how to create, modify, and delete thin provisioning profiles.

## **SAN Planner wizard: Specifying path plan**

Set multipath options for provisioning in your SAN plan.

The supported multipath driver on a host is the IBM Subsystem Device Driver (SDD). For Linux-based hosts starting with Red Hat Enterprise Linux Version 5 and SUSE Linux Enterprise Server Version 10, only the DM\_Multipath driver is available. The SDD driver is no longer supported on HP, starting with HP-UX 11i v3 with Itanium®. HP-UX 11i v3 provides a native multipathing feature which is supported by SAN Planner.

1. To specify how the driver uses the paths between the host and the storage system, select one of the following options from the **Multipath Option** field:

**Load Balancing**

Sends input/output on all paths.

**Round Robin**

Sends input/output on one path until a time interval that is set in an SDD setting at the host expires, then switches to another path.

**Fail-Over**

Sends input/output on one path until a failure occurs, then switches to another path.

On HP-UX 11i v3 with Itanium®, the multipath options used by the SAN Planner are as follows:

**Load Balancing**

If you select this option for HP-UX 11i v3, the option implemented is Least Command Load. This option selects the path with the least number of pending I/O requests for the next I/O operation.

**Round Robin**

On HP-UX 11i v3, this policy distributes the I/O load equally across all active paths irrespective of the current load on each path.

**Fail-Over**

If you select this option for HP-UX 11i v3, the policy implemented is Preferred Path. For this policy, you specify a path to an optimized controller. This path is preferred for I/O transfers. The preferred path must be set manually.

2. Select the **Specify number of paths** check box and enter the number of paths that you want to configure between the host and storage system. If you do not specify the number of paths in this field, the default is two paths (<auto>)
3. To use the paths from host to storage systems through a minimum of two fabrics, select **Use fully redundant paths (requires at least 2 fabrics)**.
4. To move to the next page in the wizard, click **Next**.

## **SAN Planner wizard: Specifying zone plan**

Use SAN Planner to specify zoning between ports on the hosts and systems in a fabric.

The supported multipath driver on a host is the IBM Subsystem Device Driver (SDD). For Linux-based hosts starting with Red Hat Enterprise Linux Version 5 and SUSE Linux Enterprise Server Version 10, only the DM\_Multipath driver is available. The SDD driver is no longer supported on HP, starting with HP-UX 11i v3 with Itanium®. HP-UX 11i v3 provides a native multipathing feature which is supported by SAN Planner.

1. To specify how the zones are created, select one of the following options from the **Automatically Create Zone** field:

**<auto-zone>**

The plan generates using the maximum number of zones without grouping the data paths based on hosts, host ports, or host bus adapters (HBAs).

**...for each host**

Creates a zone for each host.

**...for each HBA**

Creates a zone for each host bus adapter.

**...for each port**

Creates a zone for each port.

2. Optional: Select the **Specify maximum number of zones** check box and enter the maximum number of zones that you want to create between the host and storage system.
3. Optional: Select the **Maximum zone members per zone** check box and enter the maximum number of zone members for each zone.
4. Optional: If you do not want two controllers with different types to be used in the same zone, select the **No two controllers with different types should be used in the same zone** check box.
5. Optional: If you want to apply updates to the active zone set, select the **Use active zone set** check box. This action applies changes to the inactive zone set that has the same name as the active zone set and then the zone set is activated. Before using this option, ensure that the zones and zone members in the active zone set matches the zones and zone members in the inactive zone set with the same name.
6. Optional: If you want to append a zone prefix to each zone name, select the **Append zone name prefix** check box and enter a prefix.
7. To move to the next page in the wizard, click **Next**.

## **SAN Planner wizard: Selecting replication sessions**

Create a replication session for the storage being provisioned, or, if one is available, use an existing replication session.

This task is available if you have selected either **Provision storage with replication, including optional multipath and zone considerations on the source resources** or **Provision replication to existing storage** provisioning plans on the Select Planning Task page.

1. To create a replication session for the storage being provisioned, click **Create a new replication session for the storage being protected**.
2. If an existing replication session is available, click **Use an existing replication session for the storage being protected** and select the session that you want. An existing session is listed here only if the input contains a storage resource group and that storage resource group has a replication session associated with it.
3. To move to the next page in the wizard, click **Next**.

## SAN Planner wizard: Specifying replication session properties

Provide a name for the replication session and select the type of session to create.

This task is available if you have selected either **Provision storage with replication, including optional multipath and zone considerations on the source resources** or **Provision replication to existing storage** provisioning plans on the Select Planning Task page.

1. In the **Session Name** field, enter a name for the replication session.
2. Select a value from the **Session Type** list that displays those replication session types applicable to the planning you have chosen. For example, if you selected **Volumes** on the Select Provision Plan Details page, only the session types pertaining to volume planning are shown. To filter the **Session Type** list, select a value from the **Filter by Copy Technology** list. The value that you select here combined with the value in the **Location** list determines the replication sessions that are shown in the **Session Type** list. The options are:

### No filtering

No filtering is performed.

### Point-in-time Copy

The FlashCopy and Snapshot replication sessions are displayed in the **Session Type** list.

### Continuous Copy

The replication sessions that are displayed in the **Session Type** list perform continuous replication and depend on the value in the **Location** list.

3. Select a location for the target volumes or virtual disks from the **Location** list. This list is disabled if you selected **No filtering** or **Point-in-time Copy** for filtering. The options are: **Same Region**, **Across Regions**, and **Same and Across Regions**
4. Select the **Use practice volume** check box to include a practice volume in the plan. If you want to practice what you would do in the event of a disaster, without interrupting current data replication, you can use a practice volume. Practice volumes are available in Metro Mirror, Global Mirror, and Metro Global Mirror replication sessions.
5. Click the replication session type that you want to create. The copy type that is associated with the session determines the actions that can be performed against all copy sets in the session, the number of volumes or virtual disks in each copy set, and the role that each volume or virtual disk plays. The options are:
  - FlashCopy
  - VDisk Mirroring
  - Metro Mirror Failover/Failback
  - Metro Mirror Failover/Failback with Practice
  - Metro Mirror Single Direction
  - Global Mirror Either Direction with Two Site Practice
  - Global Mirror Failover/Failback
  - Global Mirror Failover/Failback with Practice
  - Global Mirror Single Direction
  - Metro Global Mirror with Practice
  - Metro Global Mirror
  - Snapshot

When you select a session type, a session image is displayed. The image is a visual aid to help you create your session. For session types other than Snapshot, this image shows the number of roles in the session and how the roles are distributed between the sites. The arrows indicate the type of copy between roles:

- A straight arrow is used to denote synchronous Metro Mirror copies.
- A zigzag arrow is used to denote asynchronous Global Copy relationships.
- A lightning bolt is used to denote FlashCopy relationships.

For Snapshot sessions, a single set of active host volumes is shown with no relationship indicators.

**Note:** For Global Mirror sessions on IBM(r) System Storage(tm) devices, SAN Planner might report a successful creation and start of the session, but the actual session start might take much longer. Verify the session by using IBM Tivoli Storage Productivity Center for Replication.

6. From the **Copy Synchronization Rate**, select the rate that data is copied from the source virtual disk to the target virtual disk. This check box is enabled only if **VDisk Mirroring** is selected in the **Session Type** list.
7. To move to the next page in the wizard, click **Next**.

### **SAN Planner wizard: Specifying a secondary location**

Specify a secondary location for creating target volumes and virtual disks. You can use an existing storage resource group (SRG), create an SRG, or use a candidate SRG. You can also select a provisioning profile for the secondary SRG.

This task is available if you have selected either **Provision storage with replication, including optional multipath and zone considerations on the source resources** or **Provision replication to existing storage** provisioning plans on the Select Planning Task page. If you did not specify an SRG in the input for the source volumes in the Specifying the Plan Content page, the secondary SRG inputs are disabled here. The list of SRGs displayed contains SRGs that are not part of any existing replication sessions. If all SRGs are part of replication sessions, these input fields are disabled.

1. To specify a secondary location, select from the following options:

#### **Secondary Candidate SRG**

Click to select the candidate SRG that you want to use. Click **View SRG** to view the candidate SRG. The secondary candidate SRG is used as the secondary source for the volume creation.

#### **Create a new storage resource group**

Click to create an SRG. Enter the name for the SRG in the **Name** field. This secondary SRG is where the replication target volume will be placed after the replication plan is executed.

#### **Use existing storage resource group**

Click to use an existing SRG. Select the SRG and click **View SRG** to view. This secondary SRG is where the replication target volume will be placed after the replication plan is executed.

2. To select a provisioning profile for the secondary SRG, select from the following options:

#### **Use the source storage resource group profile**

Click to select the provisioning profile that is associated with the source SRG.

### Use existing provisioning profile

Click to use the existing provisioning profile.

3. To move to the next page in the wizard, click **Next**.

## SAN Planner wizard: Specifying a tertiary location

Specify a tertiary location for creating target volumes and virtual disks. You can use an existing storage resource group (SRG), create an SRG, or use a candidate SRG. You can also select a provisioning profile for the tertiary SRG.

This task is available if you have selected either **Provision storage with replication, including optional multipath and zone considerations on the source resources** or **Provision replication to existing storage** provisioning plans on the Select Planning Task page, and Metro Global Mirror session as the replication session type. If you did not specify an SRG in the input for the source volumes in the Specifying the Plan Content page, the tertiary SRG inputs are disabled here. The list of SRGs displayed contains SRGs that are not part of any existing replication sessions. If all SRGs are part of replication sessions, these input fields are disabled.

1. To specify a tertiary location, select from the following options:

### Tertiary Candidate SRG

Click to select the candidate SRG that you want to use. Click **View SRG** to view the candidate SRG. The tertiary candidate SRG is used as the tertiary source for the volume creation.

### Create a new storage resource group

Click to create an SRG. Enter the name for the SRG in the **Name** field. This tertiary SRG is where the replication target volume will be placed after the replication plan is executed.

### Use existing storage resource group

Click to use an existing SRG. Select the SRG and click **View SRG** to view. This tertiary SRG is where the replication target volume will be placed after the replication plan is executed.

2. To select a provisioning profile for the tertiary SRG, select from the following options:

### Use the source storage resource group profile

Click to select the provisioning profile that is associated with the source SRG.

### Use existing provisioning profile

Click to use the existing provisioning profile.

3. To move to the next page in the wizard, click **Next**.

## SAN Planner wizard: Reviewing user selections

Review your selections in the SAN Planner wizard.

Review your selections in the wizard. If you want to change any settings, click **Back**, and make your changes. Click **Next** to generate the plan. The Getting Recommendation dialog box is displayed while the settings in the wizard are validated and the planning is performed. If the process runs without errors, the Execute Plan page is opened showing the recommended storage configuration plan. If an error occurs, a message that explains the error is displayed.

## SAN Planner wizard: Executing the plan

View the SAN Planner recommended storage configuration. Execute the plan as a job.

1. In the "Plan involves the following changes" pane, review the configuration changes that the SAN Planner job can make.
2. Click **Show Plan Topology** to look at a more detailed view of the planned storage configuration.
3. To accept the configuration changes, click one of the following options to specify when you want the job to run:

**Run now**

Click to run the job immediately.

**Run Once at**

Click to specify a date and time to run the job and select the appropriate date and time.

4. To specify the time zone that you want to use, click one of the following options:

**Use the time zone that the server runs in**

Click to use the time zone where the IBM Tivoli Storage Productivity Center server runs.

**Use this time zone**

Select the time zone from the list.

5. Click **Execute Plan** to save the plan and run the job. The job is run at the date, time, and time zone specified by the **When to run** and **How to handle time zones** settings. The plan is saved under the **SAN Planner** node in the Tivoli Storage Productivity Center navigation tree.

**Note:** For Global Mirror sessions on IBM(r) System Storage(tm) devices, SAN Planner might report a successful creation and start of the session, but the actual session start might take much longer. Verify the session using IBM Tivoli Storage Productivity Center for Replication.

## **SAN Planner wizard: Editing the plan**

Use the Edit panel to edit the plan that was created using the SAN Planner wizard or to change the run time configuration for a plan.

1. Expand the **Disk Manager** node in the navigation tree. Right-click **SAN Planner** and select a saved plan to open the wizard with the Edit Plan *TPCuserid.plan\_name* panel, where *TPCuserid* is the Tivoli Storage Productivity Center user ID and *plan\_name* is the name the plan was given when it was previously saved.
2. Click **View Job History** to view a detailed history of job runs for the plan. The Job Management panel is displayed.
3. In the **Plan involves the following changes** pane, review the configuration changes that the SAN Planner job can make.
4. Click **Show Plan Topology** to look at a more detailed view of the planned storage configuration.
5. In the **When to run pane** pane, click one of the following options to specify when you want the job to run. This pane is disabled if **Run now** is selected for the plan.

**Run now**

Click to run the job immediately.

**Run Once at**

Click to specify a date and time to run the job and select the appropriate date and time.

6. To specify the time zone that you want to use, click one of the following options in the **How to handle timezones** pane. This pane is disabled if **Run now** is selected for the plan.

**Use the time zone that the server runs in**

Click to use the time zone where the IBM Tivoli Storage Productivity Center server runs.

**Use this time zone**

Select the time zone from the list.

7. Click **Open Plan Wizard** to open the SAN Planner wizard where you can view or edit the contents of the plan.
8. Click **Save** to save the changes to the plan that you have made on this panel. This option is disabled if the **When to run** and **How to handle time zone** panes are disabled.
9. Click **Execute Plan** to run the plan immediately. This option is disabled if the **When to run** and **How to handle time zone** panes are disabled. settings.
10. Click **Cancel** to cancel the edits that you have made on this panel.

## Working with storage resource group sessions

This topic describes viewing and deleting replication sessions that are created in the SAN Planner wizard for a storage resource group.

You can delete the replication session from IBM Tivoli Storage Productivity Center and optionally also from IBM Tivoli Storage Productivity Center for Replication.

1. In the left pane, expand **Replication Manager > Storage Resource Group Sessions**. A table of the currently defined storage resource group sessions is displayed in the right pane.
2. To select a storage resource group session to delete, click a row in the table. You can use Ctrl+click and Shift+click to select multiple sessions simultaneously.
3. Click **Delete Storage Resource Group Session**. Tivoli Storage Productivity Center does not automatically delete the replication session from Tivoli Storage Productivity Center for Replication. A dialog box opens enabling you to specify how you want to delete the session. You can select from the following options:
  - Click **Yes** to remove the selected replication session from the Tivoli Storage Productivity Center database and the Tivoli Storage Productivity Center for Replication database.
  - Click **No** to remove the selected replication session from the Tivoli Storage Productivity Center database only.
  - Click **Cancel** to close the dialog box without deleting the selected replication session from the Tivoli Storage Productivity Center database or Tivoli Storage Productivity Center for Replication database.

---

## Optimizing storage configurations

This section describes how to use the Storage Optimizer to analyze storage subsystems to identify performance bottlenecks and create a recommendations report that describes improvements you can make to subsystem performance, as well as recommendations for migration and consolidation.



## Storage Optimizer overview

The Storage Optimizer uses data in the IBM Tivoli Storage Productivity Center database to analyze your storage subsystems to identify performance bottlenecks, and recommend changes to improve performance. This topic lists the supported subsystems and describes the general steps for using Storage Optimizer.

The Storage Optimizer helps you develop storage migration or storage consolidation plans, and helps you plan for the growth of your storage infrastructure. For example, you can use the Storage Optimizer to create a weekly report that provides details on the performance of your storage subsystems and makes recommendations for how to improve performance.

**Note:** The Storage Optimizer does not actually perform any migrations or make any modifications to subsystem configurations. Its primary purpose is to provide you with a performance analysis and optimization recommendations that you can choose to implement at your discretion.

**Important:** Storage Optimizer does not take into account any established replication relationships or sessions. Migration recommendations should be followed with care to ensure continuity of all replication relationships.

To use the Storage Optimizer, you must have a IBM Tivoli Storage Productivity Center Standard Edition license.

The following IBM storage subsystems or applications are supported:

- DS8000
- DS6000
- DS5000
- DS4000
- TotalStorage Enterprise Storage Server
- SAN Volume Controller
- Storwize V7000

Before running Storage Optimizer, you must set up performance monitors to collect performance data at least once a day for all storage subsystems that you want Storage Optimizer to analyze. You must also collect performance monitoring data for the back-end subsystems for SAN Volume Controller or Storwize V7000. The performance data for the subsystems produce the most accurate Storage Optimizer analysis.

The analysis is less accurate if there are any gaps in the data collection for the time interval being analyzed. It is recommended that you collect at least one week's worth of performance monitoring data before using Storage Optimizer. Providing a longer time interval for data collection will increase the accuracy of the Storage Optimizer analysis and recommendations. See "Creating storage subsystem performance monitor" on page 377 for more information.

You must provide the Storage Optimizer with the following input:

- Configuration and performance data that you collect for the supported storage subsystems before running the Storage Optimizer
- One or more selected storage subsystems that you specify as input to the Storage Optimizer.

- Connected fabric topology and zoning information will be used in the analysis if it is available.

The Storage Optimizer produces the following output:

- An analysis report that displays performance heat maps and tables that graphically illustrate the performance utilization of the storage subsystems that you specified as input.
- An optimization report that provides migration and consolidation recommendations for improving performance.

The Storage Optimizer produces more accurate results if you provide both the required and optional input data requested by the Storage Optimizer, and if you collect storage subsystem performance monitoring data over a longer time interval before using the Storage Optimizer.

**Note:** Since running the Storage Optimizer is a processor-intensive task for the IBM Tivoli Storage Productivity Center server, schedule a time to run the Storage Optimizer when processor demand on the IBM Tivoli Storage Productivity Center server is expected to be at a minimum.

The general steps for using the Storage Optimizer are as follows:

1. In the Create Analysis window, enter information in the required fields and choose **File > Save** to submit the analysis report job. See “Creating an analysis report” on page 437 for more information.
2. Review the Storage Optimizer analysis report. See “Viewing an analysis report” on page 438 for more information.
3. Create an optimization report based on this analysis. See “Creating an optimization report” on page 439 for more information.
4. View the optimization report. See “Viewing an optimization report” on page 441 for more information.
5. At your discretion, implement the recommendations suggested in the optimization report, or re-run the analysis using different inputs.

## Using the Storage Optimizer

This topic describes how to use the Storage Optimizer to analyze storage subsystem performance and create an optimization report that includes recommendations to improve performance.

Before running Storage Optimizer, you must set up performance monitors and collect performance monitoring data for all storage subsystems that you want Storage Optimizer to analyze. Storage Optimizer analyzes performance monitoring data for the time interval you specify. The analysis will be less accurate if there are any gaps in the data collection for the time interval being analyzed. For best results, collect at least one week's worth of performance monitoring data before using Storage Optimizer. Providing a longer time interval for data collection will increase the accuracy of the Storage Optimizer analysis and recommendations. See “Creating storage subsystem performance monitor” on page 377 for more information.

**Tip:** The Storage Optimizer does not make any modifications to subsystem configurations. Its primary purpose is to provide you with a performance analysis and optimization recommendations that you can choose to implement at your discretion.

## Creating an analysis report

Creating a performance analysis report is the first step in using Storage Optimizer.

When you specify a start time for using performance monitor data in an optimization plan, the start time might not coincide with the start of a performance monitor data collection interval. If the requested start time is after the start of a data collection interval and the duration of the data collection interval is long enough so that it ends at or after the requested start time, Tivoli Storage Productivity Center includes all data from the collection interval that began before the requested start time. This might include data collected before the start time that you specified. This processing ensures that the maximum amount of available performance data is utilized. For example, if a performance monitor collects data every 20 minutes starting at 01:00, but you request performance data starting at 1:15, Tivoli Storage Productivity Center still analyzes performance data starting at 01:00 that has a duration of 20 minutes.

To create the analysis, complete the following steps:

1. Navigate to one of the following nodes:
  - **IBM Tivoli Storage Productivity Center > Analytics > Storage Optimizer**
  - **Disk Manager > Storage Optimizer.**
2. Right-click **Storage Optimizer** and select **Create Analysis** to open the Create Analysis panel. Use this panel to specify the storage subsystems that you want to analyze.
3. Optional: In the **Create Analysis** panel, enter a short description of the analysis job in the **Description** field.
4. The **Select Storage** pane lists the storage subsystems that you want to include in the analysis. Use the **Add** button to add elements to the plan:
  - a. Click **Add**. The **Optimizer Selection** panel opens. This provides a topology view of the current system configuration. Select the storage subsystems that you want to include in the analysis.

**Tip:** The Topology Viewer pane displays the configuration from which selections are made. The Select Elements pane displays the selections that will be used in the analysis.
  - Note:** When using the Select Elements pane, the subsystem selections do not need to be in the same fabric.
  - b. Double-click the Storage box title to expand the storage view into the L0 level view of the available subsystems.
  - c. Click an element to select it. To select multiple elements, press and hold the **CTRL** key and click each element icon.
  - d. Click **>>** to move the selected elements into the **Select Elements** pane. If you decide to remove the selected elements from the **Select Elements** pane, click **<<**.
  - e. When you are satisfied with all your selections, click **OK**. You return to the Storage Optimizer Create Analysis panel with the selections displayed in the **Select Storage** pane.
5. Optional: select an element to remove it by clicking **Remove**.
6. In the **Performance Time Interval** section, for **Data Type**, choose Daily to get average performance data for each 24 hour period, or choose Sample/Hourly to analyze a specific time period each day. Select the **Start Date** and **End Date** to define the date interval that the Storage Optimizer uses to extract performance monitoring data from the IBM Tivoli Storage Productivity Center

database. Performance monitoring data must exist in the database for the entire interval you specify. The default date interval is the start and end date of the current date. If you choose Sample/Hourly, specify the **Start Time** and **End Time** of the daily time period that you want to analyze. For example, if you choose Sample/Hourly with November 1 as the start date, November 15 as the end date, and 6 PM as the start time and 8 PM as the end time, then data will be analyzed between 6 PM and 8 PM only for days between November 1 and November 15.

7. Click **Run Now** to begin the analysis as soon as you save the analysis job using **File > Save**. Alternatively, click **Run Once at** to begin the analysis at a time you specify. When you save the analysis job using **File > Save**, the analysis job appears under the **Storage Optimizer** node in the navigation tree, but the job status is not displayed until the job has started running.
8. For **How to handle time zones**, specify which time zone to use.
9. Click **File > Save** to save and submit the analysis report job. After the analysis report job starts, you can view the analysis job status under the **Storage Optimizer** node in the navigation tree.
10. When the analysis report job is complete, select it to view the analysis report.

### Viewing an analysis report

The analysis report displays a performance analysis for storage systems, and lets you generate an optimization report that lists storage migration and consolidation recommendations. You can experiment with different migration and consolidation scenarios to help you achieve the desired performance improvements for your storage infrastructure.

To view the analysis report:

1. Navigate to one of the following nodes:
  - **IBM Tivoli Storage Productivity Center > Analytics > Storage Optimizer**
  - **Disk Manager > Storage Optimizer**.
2. Expand the **Storage Optimizer** node and select a completed (green) analysis job to display the analysis report. By default, the heat maps display the performance utilization of all components.
3. To display the performance utilization for a specific component, select the component from the **Heat map based on** list.
4. Drag the **Performance Threshold** slider to change the performance threshold for all the selected subsystems. The heat maps are automatically updated to reflect changes in the performance threshold. This allows you to see the performance threshold at which different subsystems or components become bottlenecked.

There is one heat map for each storage subsystem included in the analysis. The heat map uses different colors to represent the actual performance of a storage subsystem, as measured against the default performance threshold of 80%. If you change the performance threshold, the heat maps are automatically updated to display the changes.

Each cell represents a storage pool. If there aren't enough storage pools to fill an entire row, the remaining cells in that row will be gray. Select a cell to display more information about that storage pool. Right-click a cell and choose **show optimizer view** to display performance metrics for volumes in the selected pool.

Green, blue, yellow, orange, and red cells indicate storage pool performance as follows:

**Green** storage pools that are performing at less than or equal to 25% of the performance threshold.

**Blue** storage pools that are performing at less than or equal to 50% of the performance threshold.

**Yellow** storage pools that are performing at less than or equal to 75% of the performance threshold.

**Orange** storage pools that are performing at less than or equal to 100% of the performance threshold.

**Red** storage pools that exceed the performance threshold.

The Performance tables represent the same information that is displayed in each heat map, but in more detail. There is one table for each storage subsystem. Each row represents a storage pool. You can sort the table by clicking a column. Column descriptions are as follows:

**Pool Name**  
the name of the storage pool.

**Utilization**  
the aggregated performance utilization of all components, as measured against the performance threshold.

**Host Adaptor**  
the performance utilization of the host adaptor ports on the storage subsystem.

**Hard Disk**  
the performance utilization of the disk arrays.

**Controller**  
processor and memory performance utilization.

**Device Adapter**  
the performance utilization of the device adapter which connects the controller to the disk arrays.

**Space** The physical disk space utilization.

## Creating an optimization report

You can create an optimization report that lists storage migration and consolidation recommendations. You can experiment with different migration and consolidation scenarios to help you achieve the desired performance improvements for your storage infrastructure.

To create the optimization report, follow these steps:

1. Navigate to one of the following nodes:
  - **IBM Tivoli Storage Productivity Center > Analytics > Storage Optimizer**
  - **Disk ManagerStorage Optimizer.**
2. Select a completed analysis job to display the analysis report.
3. To include a storage pool or storage subsystem as a source entity that you want to migrate or consolidate to improve performance:
  - a. In the heat map or table, select a storage pool or storage subsystem. Select **Select all pools for storage subsystem** to select all storage pools in a storage subsystem so that you can include them as source entities.

- b. Click >> to the left of **Source Entities** to add the selected item to the source entities. The same storage pool cannot be added as both a source and target entity.
- c. For the list of source entities, if you select **Make retirement recommendations for the selected entity**, the optimization report will make retirement recommendations for all source entities.
- d. To remove an item from the list of source entities, select it and click << to the left of **Source Entities**.


**Note:** The values in the optimization report are estimates and could be influenced by data activity on the selected entity

The Storage Optimizer does not support Easy Tier (mixed pools) as source in a migration or consolidation plan. However, Easy Tier pools are displayed in the Performance Heat Maps and Performance Tables when you create an optimization report. These pools are marked with the text **E.T.** and cannot be selected.

4. To include a storage pool or storage subsystem as a target entity that you want to keep after storage migration and consolidation:
  - a. In the heat map or table, select a storage pool or storage subsystem. Select **Select all pools for storage subsystem** to select all storage pools in a storage subsystem so that you can include them as target entities.
  - b. Click >> to the left of **Target Entities** to add the selected item to the target entities. The same storage pool cannot be added as both a source and target entity.
  - c. To remove an item from the list of target entities, select it and click << to the left of **Target Entities**.

When the Storage Optimizer considers potential pools as targets of a migration, it takes into consideration whether the source and target pools have the same RAID level, format (CKD or FB), and whether or not the potential target has enough available capacity. The Storage Optimizer does not support Easy Tier (mixed pools) as target entities in a migration or consolidation plan.

5. Choose when to run the optimization report using one of the following options:
  - **Run Now:** creates the optimization report as soon as you save your report settings by clicking **File > Save**. After the report job begins, you can view the report and job status under the **View Previously Run Optimization Reports** pane.
  - **Run Once at:** begins the report job at a specified date and time. When you save the report job by clicking **File > Save**, the report job is displayed under the **View Previously Run Optimization Reports** pane, but the job status will not appear until the job has started running.
6. Drag the **Performance Threshold** slider to change the performance threshold for the selected subsystems and pools. The heat maps are automatically updated to reflect changes in the performance threshold. This allows you to see the performance threshold at which different subsystems and pools become bottlenecked. Keep in mind that if you select the **Make retirement recommendations for the selected entity** option, the threshold you choose using the **Performance Threshold** slider determines if the move to the target subsystem will keep the target subsystem below the chosen threshold.
7. Click **File > Save** to save and submit the optimization report job.
8. To update the job status for all optimization reports, click **Refresh Job Status**.

9. To view a completed optimization report, click  .
10. To see how the optimization recommendations change if you use a different performance threshold, select a different performance threshold and run another optimization report. You can continue creating optimization reports using different performance thresholds until you achieve the expected performance improvements.


## Viewing an optimization report

The optimization report displays a storage migration and consolidation report that lists recommendations for improving storage subsystem performance.

The Storage Optimizer does not actually perform any migrations or make any modifications to subsystem configurations. Its primary purpose is to provide you with recommendations that you can choose to implement at your discretion.

**Important:** Storage Optimizer does not take into account any established replication relationships or sessions. Migration recommendations should be followed with care to ensure continuity of all replication relationships.

To view the optimization report:

1. Navigate to one of the following nodes:
  - **IBM Tivoli Storage Productivity Center > Analytics > Storage Optimizer**
  - **Disk ManagerStorage Optimizer.**
2. Select a completed analysis job.
3. Optional: Under the **View Previously Run Optimization Reports** pane, click **Refresh Job Status** to update the job status for all optimization reports.
4. Click  to open an optimization report.
5. To view the results of implementing only some of the recommendations, remove the check mark next to **Select all recommendations** and select a subset of the recommendations from the table of recommendations. Only the selected recommendations will be included in the printed report.
6. To display the performance utilization for a specific component, select the component next to **Heat map based on**.
7. To print the report, click the **Print** button in the optimization report. If a PDF printer driver is installed, you can print the report to a PDF file. Alternatively, you can print the report to a printer. The printed report includes an explanation of each recommendation, and sample migration scripts for the SAN Volume Controller.

The optimization recommendations include the following information:

### Recommendation number

The sequential number associated with the recommendation.

### Source subsystem

The name of the source subsystem.

### Target subsystem

The name of the target subsystem.

### Source volume

The name of the source volume.

**Target volume**

The name of the target volume.

**Source pool**

The name of the source pool.

**Target pool**

The name of the target pool.

**Reason**

The reason why the storage pool utilization has exceeded the performance threshold.

**Port information**

The recommendation on which ports need to be configured on the target subsystem for the target volume. This information only appears in the printed report.

**Zone information**

The recommended zoning changes necessary so that the target volume is visible to the host. This information only appears in the printed report.

**SVC pseudoscript**

The sample migration script that serves as a guide to the commands that you must enter using the SAN Volume Controller command-line interface. This information only appears in the printed report.

After you implement the recommendations, if you want to verify that you have achieved the expected performance improvements, you must first collect more performance monitoring data before running the Storage Optimizer again.

**Deleting analysis jobs**

You can delete analysis job definitions or individual analysis jobs if you no longer need them.

To delete an analysis job definition or job run:

1. Navigate to one of the following nodes:
  - **IBM Tivoli Storage Productivity Center > Analytics > Storage Optimizer**
  - **Disk ManagerStorage Optimizer> Disk Manager.**
2. Expand the **Storage Optimizer** node.
3. Right-click an analysis job definition or an analysis job.
4. Click **Delete**.

Keep in mind that deleting an analysis job definition also deletes all analysis jobs and optimization reports that are associated with that job definition. If you delete an analysis job, all optimization reports associated with that job are also deleted.

**Deleting an optimization report**

You can delete optimization reports if you no longer need them.

To delete an optimization report job:

1. Navigate to one of the following nodes:
  - **IBM Tivoli Storage Productivity Center > Analytics > Storage Optimizer**
  - **Disk ManagerStorage Optimizer.**
2. Expand the **Storage Optimizer** node.



3. Select a completed analysis job to display the list of optimization reports that are associated with that analysis job.
4. Under **View Previously Run Optimization Reports**, right-click a completed report and choose **Delete**.

### Retrieving an optimization report job definition

You can retrieve the report job definition that was used to generate an optimization report.

To retrieve the optimization report job definition associated with an optimization report:

1. Navigate to one of the following nodes:
  - **IBM Tivoli Storage Productivity Center > Analytics > Storage Optimizer**
  - **Disk ManagerStorage Optimizer**.
2. Expand the **Storage Optimizer** node.
3. Select a completed analysis job to display the list of optimization reports that are associated with that analysis job under **View Previously Run Optimization Reports**.
4. Right-click an optimization report and choose **Retrieve**.

The **Generate Optimization Report** pane displays the inputs that were used to create the selected optimization report.

### Understanding zoning recommendations

This topic describes the zoning recommendations provided by Storage Optimizer in the printed optimization report.

**Note:** Storage Optimizer only provides zoning recommendations in the printed optimization report. The report lists the zoning recommendations for all volumes under each volume.

Where possible, Storage Optimizer provides zoning recommendations for each migration recommendation. The zoning recommendations will be to put the target subsystem's port(s) into the same zones as that of the source subsystem's ports, regardless of which zones the host is also a member of.

For example, suppose you want Storage Optimizer to provide zoning recommendations using the following source and target subsystems.

Source subsystem A includes the following:

- Port 1 (member of zones: host1\_zone, host2\_zone, host3\_zone)
- Port 2 (member of zones: host4\_zone, host5\_zone, host6\_zone)
- Volume X (mapped to host1 using Port 1)
- Volume Y (mapped to host4 using Port 2)

Target subsystem B includes the following:

- Port 1
- Port 2

Suppose that Storage Optimizer recommends migrating Volume X from Source subsystem A to Target subsystem B. The zoning recommendations that accompany this migration recommendation will be to add Target subsystem B's Port 1 to zones host1\_zone, host2\_zone, and host3\_zone.

## Enforcing storage policies

Enforcing policies for storage usage can be critical to ensuring the proper and most cost-effective use of storage devices. Policies are enforced when conditions that you defined are detected by a monitoring job. Use quotas, constraints, filesystem extensions, and scheduled actions to manage and enforce your storage policies.

Use the **Policy Management** node in Data Manager and Data Manager for Databases to manage and enforce storage policies.

For example, if you want to enforce a policy that limits the amount of space a group of users can consume on the network, use a Data Manager quota. That quota will be triggered: 1) after a scan is run and 2) the scan detected that a group of users is consuming more storage space than was indicated in the quota.

The following table describes the jobs you can use to enforce policies for storage usage in your environment:

Table 22. Descriptions of policy types and their triggering conditions

| Policy Type / Manager                                                                                                                                                                                                                                                                                                                                   | What can I do with this job?                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | What conditions can trigger an alert?                                                                                                                                                                                                                                       |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Quotas</b><br>(Data Manager, Data Manager for Databases) <ul style="list-style-type: none"> <li>• Policy Management &gt; Quotas &gt; User</li> <li>• Policy Management &gt; Quotas &gt; OS User Group</li> </ul>                                                                                                                                     | Use quotas to define limits on and monitor the amount of storage that a user or a group of users can consume. You can set quotas to specify limits for storage usage at different levels, which enables you to control precisely how much storage a user can consume on specific storage entities: <ul style="list-style-type: none"> <li>• Define quotas in Data Manager to set limits on the amount of storage that a user or a group of users can consume at a network, filesystem, or computer level.</li> <li>• Define quotas in Data Manager for Databases to set limits on the amount of storage that a user or a group of users can consume at a network, instance, database, or tablespace level.</li> </ul> Monitoring job that detects conditions in quotas: scan | <b>User consumes more than <math>n</math></b> , where $n$ represents a number measured in kilobytes, megabytes, or gigabytes.<br><br><b>OS User Group consumes more than <math>n</math></b> , where $n$ represents a number measured in kilobytes, megabytes, or gigabytes. |
| <b>Network Appliance Quotas</b><br>(Data Manager) <ul style="list-style-type: none"> <li>• Policy Management &gt; Network Appliance Quotas &gt; Schedules</li> <li>• Policy Management &gt; Network Appliance Quotas &gt; Imported User Quotas</li> <li>• Policy Management &gt; Network Appliance Quotas &gt; Imported OS User Group Quotas</li> </ul> | Use Network Appliance Quotas to import, view, and report on the NetApp Quotas defined on NAS filers.<br><br>Monitoring job that detects conditions in Network Appliance quotas: scan                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>% of hard limit</b> , where $n$ represents a percentage of the hard limit defined in the NetApp Quota.                                                                                                                                                                   |

Table 22. Descriptions of policy types and their triggering conditions (continued)

| Policy Type / Manager                                                                                                                              | What can I do with this job?                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | What conditions can trigger an alert?                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>Constraints</b><br/>(Data Manager)</p> <p><b>Policy Management &gt; Constraints</b></p>                                                      | <p>Use constraints to:</p> <ul style="list-style-type: none"> <li>define the acceptable and unacceptable file types, file sizes, and file owners for computers in your environment. For example, you can use constraints to alert you when users store certain files (such as MP3 or .avi files) on a monitored computer.</li> <li>Request an IBM Tivoli Storage Productivity Center archive and backup of the largest violating files identified by a constraint. IBM Tivoli Storage Manager protects your organization's data from hardware failures and other errors by storing backup and archive copies of data on offline storage.</li> </ul> <p>Monitoring job that detects conditions in constraints: scan</p> | <p><b>Violating files consume more than n</b>, where <i>n</i> represents a number measured in kilobytes, megabytes, or gigabytes.</p>                                                                                                                                                                                                                                                                                                                                 |
| <p><b>Filesystem Extension</b><br/>(Data Manager)</p> <p><b>Policy Management &gt; Filesystem Extension</b></p>                                    | <p>Use file system extension to create additional space in the local file systems of managed hosts. You can extend file systems manually, or set up policy to do it automatically. You can extend file systems at a specified time, or when utilization reaches a specified threshold.</p> <p>Monitoring job that detects conditions in file system extensions: scan</p>                                                                                                                                                                                                                                                                                                                                               | <p><b>A file system extension action started automatically</b></p>                                                                                                                                                                                                                                                                                                                                                                                                    |
| <p><b>Scheduled Actions</b><br/>(Data Manager, Data Manager for Databases)</p> <p><b>Policy Management &gt; Scheduled Actions &gt; Scripts</b></p> | <p>Use scheduled actions to schedule scripts to run against computers and computer groups. Scheduled actions are script-based, enabling you to use any third-party tools for actions, such as recovery or provisioning.</p> <p>Monitoring job that detects conditions in scheduled actions: n/a</p>                                                                                                                                                                                                                                                                                                                                                                                                                    | <p>Not applicable. The Scheduled Action facility enables you to run scripts according to the schedule that <i>you define</i>. You do not have to associate the script with an alert.</p>                                                                                                                                                                                                                                                                              |
| <p><b>Archive/Backup</b><br/>(Data Manager)</p> <p><b>Policy Management &gt; Archive/Backup</b></p>                                                | <p>View and edit archive and backup jobs that were created based on files selected from reports in the Reporting facility.</p> <p>Monitoring job that detects conditions in archive/backups: scan</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <p>Not applicable. Any archive/backup jobs displayed under this node were created based on files shown in the following reports:</p> <ul style="list-style-type: none"> <li>Usage &gt; Files &gt; Largest Files</li> <li>Usage &gt; Files &gt; Most Obsolete Files</li> <li>Usage &gt; Files &gt; Orphan Files</li> <li>Usage &gt; Files &gt; Duplicate Files</li> <li>Backup &gt; Most At Risk Files</li> <li>Usage Violations &gt; Constraint Violations</li> </ul> |

## Setting storage usage quotas

Use quotas to set limits on the amount of storage space that a user or group of users can consume. You can use Data Manager and Data Manager for Databases to set quotas at the filesystem, computer, network, database, tablespace, and instance levels.

Use quotas in the to set limits on the amount of storage that a user or a group of users can consume.

The following steps provide an overview of how quotas work:

1. A scan gathers extensive information about the storage resources within an enterprise. This information is stored in the database repository.
2. When a quota job runs, it checks the information in the repository to determine if any quota violations exist.
3. If there are any violations, an alert is triggered and the appropriate action is taken.

**Note:** Because the information analyzed by quotas is provided by scans, you should schedule a quota to run after the scan to collect that information. This ensures that quotas are using the latest information about an enterprise to determine violations.

You can define two different kinds of quotas: User and OS User Group.

- **User Quota.** Select users and groups of users for a quota that have been defined within Data Manager or Data Manager for Databases.
- **OS User Group Quota.** Select OS User Groups for a quota. OS User Groups are groups of users that have been defined in the operating system outside of Data Manager. You can create Groups that contain many OS User Groups. When creating a quota of this type, you can select individual OS User Groups and groups that consist of multiple OS User Groups.

By providing you with the ability to define quotas at different levels, you can control precisely how much storage a user or user group can consume on specific filesystems, computers, network, databases, tablespaces, and instances:



| Level                      | Description                                                                                                                                                                                                                                                                                                              |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Filesystem</b>          | Set storage usage limits on specific filesystems and groups of filesystems that are monitored by Data Manager. This means that you will be alerted if a user or group of users has exceeded the storage usage limit you defined for the specific filesystem or groups of filesystems in the quota.                       |
| <b>Computer</b>            | Set storage usage limits on specific computers and groups of computers that are monitored by Data Manager. This means that you will be alerted if a user or group of users has exceeded the storage usage limit you defined for the specific computers or groups of computers in the quota.                              |
| <b>Database-Tablespace</b> | Set storage usage limits on specific databases, tablespaces, groups of databases, and groups of tablespaces that are monitored by Data Manager for Databases. This means that you will be alerted if a user or group of users has exceeded the storage usage limit you defined for those storage resources in the quota. |

| Level    | Description                                                                                                                                                                                                                                                                                                                        |
|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Instance | Set storage usage limits on specific instances, computers, and groups of computers that are monitored by Data Manager for Databases. This means that you will be alerted if a user or group of users has exceeded the storage usage limit you defined for the specific instances, computers, and groups of computers in the quota. |
| Network  | Set network-level storage usage limits for <i>multiple</i> computers, groups of computers, filesystems, databases, and tablespaces within your network that are monitored by Data Manager and Data Manager for Databases.                                                                                                          |

## Creating a network-wide quota

Learn how to create a quota that generates an alert when a user or group of users consumes more than a specified amount of storage space across the filesystems and computers on a network.

To create a network-wide quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager > Policy Management > Quotas**.
2. Expand the **User** or **OS User Group** node depending upon which entity you want to apply a quota.
3. Select the level at which you want to apply the quota by right-clicking **Network**.
4. Select **Create Quota** from the pop-up menu. The Create Quota window opens.
5. In the **Description** field, type a brief description of the quota.
6. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP\_BY\_OWNER type profiles shown. Make sure that the selected profile is used in scans of all the filesystems chosen on the **Filesystems** tab. This ensures that the quota is fairly enforced for all users.
7. Select users or user groups against which you want to apply the quota:
  - a. In the **Available** field, click the users or groups of users against which you want to apply the quota.
  - b. Click . The users or groups are moved to the **Current Selections** field.
8. Click the **Filesystems** tab; the Filesystems page opens.
9. Select filesystems, filesystem groups, computers, and computers groups on which you want to limit the storage usage of users and user groups.
  - a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups displayed in this field were created using **Data Manager > Monitoring > Groups**.
  - b. Click . The resources are moved to the **Current Selections** field.
10. Click the **When to Run** tab; the When to Run page opens.
11. In the **How often to run** area, schedule how often you want the quota to be checked. You can choose from the following actions:

### Run Now

Check the quota immediately.

### Run Once at

Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

**Note:** The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

### Run Repeatedly

Check the quota repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
  - Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
  - Indicate which days on which run the job. A check mark is displayed next to the days you select.
12. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
- a. 9:00 AM of the time zone where the Data server resides
  - b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

### Use the timezone that the server runs in

Select this option to use the time zone of the location where the Data server resides.

### Use this timezone

Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field is activated.

13. Click the **Alert** tab; the Alert page opens.
14. In the **Triggering Condition** area, specify the event that will trigger the alert:
- a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where *n* represents a value that you define in the **Value** field. No other conditions are available.
  - b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota is triggered.
  - c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.
15. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition. You can choose from the following actions:

### SNMP trap

An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

### **TEC / OMNIbus Event**

A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, and so on.

### **Login notification**

The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

### **Login ID**

The ID of the user who receives alert notification upon logging in to the system.

### **Windows Event Log, UNIX Syslog**

The ID of the user who receives alert notification upon logging in to the system.

### **Event Type (Windows Event Log only)**

The type of event that is recorded to the OS log.

### **Facility (UNIX Syslog only)**

The location where the UNIX Syslog is sent. You can select User or Local.

### **Level (UNIX Syslog only)**

The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

### **Run script**

A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window are passed to the script when the alert is triggered.

**E-mail** An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit e-mail** to change the actual email message.



16. Save the quota:
  - a. Click **File > Save**. The Save As window opens.
  - b. In the **Specify Quota name** field, type a name for the quota and click **OK**. The quota is displayed in the navigation tree pane.

## **Creating a network-wide quota for databases**

Learn how to create a quota that generates an alert when a user or group of users consumes more than a specified amount of RDBMS storage space across the databases, tablespaces, groups of database-tablespaces, computers, and groups of computers on a network.

To create a network-wide quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager for Databases > Policy Management > Quotas**.
2. Select the level at which you want to apply the quota by right-clicking **Network**.
3. Select **Create Quota** from the pop-up menu. The Create Quota window opens.

4. In the **Description** field, type a brief description of the quota.
5. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP\_BY\_OWNER type profiles shown. Make sure that the selected profile is used in scans of all the databases and tablespaces chosen on the **Databases-Tablespaces** tab. This ensures that the quota is fairly enforced for all users.
6. Select users or user groups against which you want to apply the quota:
  - a. In the **Available** field, click the users or groups of users against which you want to apply the quota.
  - b. Click . The users or groups are moved to the **Current Selections** field.
7. Click the **Databases-Tablespaces** tab; the Databases-Tablespaces page opens.
8. Select the databases, tablespaces, database-tablespaces groups, computers, and computer groups on which you want to limit the storage usage of users and user groups.
  - a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups displayed in this field were created using **Data Manager for Databases > Monitoring > Groups**.
  - b. Click . The resources are moved to the **Current Selections** field.
9. Click the **When to Run** tab; the When to Run page opens.
10. In the **How often to run** area, schedule how often you want the quota to be checked. You can choose from the following actions:

**Run Now**

Check the quota immediately.

**Run Once at**

Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

**Note:** The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

**Run Repeatedly**

Check the quota repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
- Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark is displayed next to the days you select.

11. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
  - a. 9:00 AM of the time zone where the Data server resides



- b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

**Use the timezone that the server runs in**

Select this option to use the time zone of the location where the Data server resides.

**Use this timezone**

Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field is activated.

- 12. Click the **Alert** tab; the Alert page opens.
- 13. In the **Triggering Condition** area, specify the event that will trigger the alert:
  - a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where *n* represents a value that you define in the **Value** field. No other conditions are available.
  - b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota is triggered.
  - c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.
- 14. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition. You can choose from the following actions:

**SNMP trap**

An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

**TEC / OMNIbus Event**

A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, and so on.

**Login notification**

The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

**Login ID**

The ID of the user who receives alert notification upon logging in to the system.

**Windows Event Log, UNIX Syslog**

The ID of the user who receives alert notification upon logging in to the system.

**Event Type (Windows Event Log only)**

The type of event that is recorded to the OS log.

**Facility (UNIX Syslog only)**

The location where the UNIX Syslog is sent. You can select User or Local.

### Level (UNIX Syslog only)

The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

### Run script

A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window are passed to the script when the alert is triggered



**E-mail** An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit e-mail** to change the actual email message.

15. Save the quota:
  - a. Click **File > Save**. The Save As window opens.
  - b. In the **Specify Quota name** field, type a name for the quota and click **OK**. The quota is displayed in the navigation tree pane

## Creating a computer quota

Learn how to create a quota that generates an alert when a user or group of users consumes more than a specified amount of storage space across computers and groups of computers.

To create a computer quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager > Policy Management > Quotas**.
2. In the **Description** field, type a brief description of the job.
3. Expand the **User** or **OS User Group** node depending upon which entity you want to apply a quota.
4. Select the level at which you want to apply the quota by right-clicking **Computer**.
5. Select **Create Quota** from the pop-up menu. The Create Quota window opens.
6. In the **Description** field, type a brief description of the quota.
7. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP\_BY\_OWNER type profiles shown.
8. Select users or user groups against which you want to apply the quota:
  - a. In the **Available** field, click the users or groups of users against which you want to apply the quota.
  - b. Click . The users or groups are moved to the **Current Selections** field.
9. Click the **Computers** tab; the Computers page opens.
10. Select computers and computers groups on which you want to limit the storage usage of users and user groups.
  - a. In the **Available** field, click the computers and computers groups to which you want to apply the quota. The groups displayed in this field were created using **Data Manager > Monitoring > Groups**.
  - b. Click . The resources are moved to the **Current Selections** field.
11. Click the **When to Run** tab; the When to Run page opens.
12. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

**Run Now**

Check the quota immediately.

**Run Once at**

Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

**Note:** The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

**Run Repeatedly**

Check the quota repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
- Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark is displayed next to the days you select.

13. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
  - a. 9:00 AM of the time zone where the Data server resides
  - b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

**Use the timezone that the server runs in**

Select this option to use the time zone of the location where the Data server resides.

**Use this timezone**

Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field is activated.

14. Click the **Alert** tab; the Alert page opens.
15. In the **Triggering Condition** area, specify the event that will trigger the alert:
  - a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where *n* represents a value that you define in the **Value** field. No other conditions are available.
  - b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota is triggered.
  - c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.
16. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition. You can choose from the following actions:

**SNMP trap**

An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

**TEC / OMNIbus Event**

A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, and so on.

**Login notification**

The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

**Login ID**

The ID of the user who receives alert notification upon logging in to the system.

**Windows Event Log, UNIX Syslog**

The ID of the user who receives alert notification upon logging in to the system.

**Event Type (Windows Event Log only)**

The type of event that is recorded to the OS log.

**Facility (UNIX Syslog only)**

The location where the UNIX Syslog is sent. You can select User or Local.

**Level (UNIX Syslog only)**

The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

**Run script**

A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window are passed to the script when the alert is triggered.

**E-mail** An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit e-mail** to change the actual email message.



17. Save the quota:
  - a. Click **File > Save**. The Save As window opens.
  - b. In the **Specify Quota name** field, type a name for the quota and click **OK**. The quota is displayed in the navigation tree pane.

**Creating a filesystem quota**

Learn how to create a quota that generates an alert when a user or group of users consumes more than a specified amount of storage space across filesystems, groups of filesystems, computers, and groups of computers.

To create a filesystem quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager > Policy Management > Quotas**.

2. Expand the **User** or **OS User Group** node depending upon which entity you want to apply a quota.
3. Select the level at which you want to apply the quota by right-clicking **Filesystem**.
4. Select **Create Quota** from the pop-up menu. The Create Quota window opens.
5. In the **Description** field, type a brief description of the job.
6. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP\_BY\_OWNER type profiles shown. Make sure that the selected profile is used in scans of all the filesystems chosen on the **Filesystems** tab. This ensures that the quota is fairly enforced for all users.
7. Select users or user groups against which you want to apply the quota:
  - a. In the **Available** field, click the users or groups of users against which you want to apply the quota.
  - b. Click . The users or groups are moved to the **Current Selections** field.
8. Click the **Filesystems** tab; the Filesystems page opens.
9. Select filesystems, filesystems, computers, and computer groups on which you want to limit the storage usage of users and user groups.
  - a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups displayed in this field were created using **Data Manager > Monitoring > Groups**.
  - b. Click . The resources are moved to the **Current Selections** field.
10. Click the **When to Run** tab; the When to Run page opens.
11. In the **How often to run** area, schedule how often you want the quota to be checked. You can choose from the following actions:

#### **Run Now**

Check the quota immediately.

#### **Run Once at**

Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

**Note:** The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

#### **Run Repeatedly**

Check the quota repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
- Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark is displayed next to the days you select.

12. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
  - a. 9:00 AM of the time zone where the Data server resides
  - b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

**Use the timezone that the server runs in**

Select this option to use the time zone of the location where the Data server resides.

**Use this timezone**

Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field is activated.

13. Click the **Alert** tab; the Alert page opens.
14. In the **Triggering Condition** area, specify the event that will trigger the alert:
  - a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where *n* represents a value that you define in the **Value** field. No other conditions are available.
  - b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota is triggered.
  - c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.
15. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition. You can choose from the following actions:

**SNMP trap**

An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

**TEC / OMNIBUS Event**

A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, and so on.

**Login notification**

The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

**Login ID**

The ID of the user who receives alert notification upon logging in to the system.

**Windows Event Log, UNIX Syslog**

The ID of the user who receives alert notification upon logging in to the system.

**Event Type (Windows Event Log only)**

The type of event that is recorded to the OS log.

**Facility (UNIX Syslog only)**

The location where the UNIX Syslog is sent. You can select User or Local.

**Level (UNIX Syslog only)**

The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

**Run script**

A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window are passed to the script when the alert is triggered



**E-mail** An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit e-mail** to change the actual email message.

16. Save the quota:
  - a. Click **File > Save**. The Save As window opens.
  - b. In the **Specify Quota name** field, type a name for the quota and click **OK**. The quota is displayed in the navigation tree pane

**Creating an instance quota**

Learn how to create a quota that generates an alert when a user or group of users consumes more than a specified amount of RDBMS storage space across instances, computers, and groups of computers.

To create an instance quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager for Databases > Policy Management > Quotas**.
2. Select the level at which you want to apply the quota by right-clicking **Instance**.
3. Select **Create Quota** from the pop-up menu. The Create Quota window opens.
4. In the **Description** field, type a brief description of the quota.
5. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP\_BY\_OWNER type profiles shown.
6. Select users or user groups against which you want to apply the quota:
  - a. In the **Available** field, click the users or groups of users against which you want to apply the quota.
  - b. Click . The users or groups are moved to the **Current Selections** field.
7. Click the **Instances** tab; the Instances page opens.
8. Select the instances, computers, and computer groups on which you want to limit the storage usage of users and user groups.
  - a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups displayed in this field were created using **Data Manager for Databases > Monitoring > Groups**.
  - b. Click . The resources are moved to the **Current Selections** field.
9. Click the **When to Run** tab; the When to Run page opens.
10. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

### Run Now

Check the quota immediately.

### Run Once at

Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

**Note:** The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

### Run Repeatedly

Check the quota repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
- Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark is displayed next to the days you select.

11. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
  - a. 9:00 AM of the time zone where the Data server resides
  - b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

### Use the timezone that the server runs in

Select this option to use the time zone of the location where the Data server resides.

### Use this timezone

Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field is activated.

12. Click the **Alert** tab; the Alert page opens.
13. In the **Triggering Condition** area, specify the event that will trigger the alert:
  - a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where *n* represents a value that you define in the **Value** field. No other conditions are available.
  - b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota is triggered.
  - c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.
14. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition. You can choose from the following actions:



**SNMP trap**

An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

**TEC / OMNIbus Event**

A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, and so on.

**Login notification**

The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

**Login ID**

The ID of the user who receives alert notification upon logging in to the system.

**Windows Event Log, UNIX Syslog**

The ID of the user who receives alert notification upon logging in to the system.

**Event Type (Windows Event Log only)**

The type of event that is recorded to the OS log.

**Facility (UNIX Syslog only)**

The location where the UNIX Syslog is sent. You can select User or Local.

**Level (UNIX Syslog only)**

The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

**Run script**

A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window are passed to the script when the alert is triggered.

**E-mail** An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit e-mail** to change the actual email message.

15. Save the quota:



- a. Click **File > Save**. The Save As window opens.
- b. In the **Specify Quota name** field, type a name for the quota and click **OK**. The quota is displayed in the navigation tree pane.

**Creating a database-tablespace quota**

Learn how to create a quota that generates an alert when a user or group of users consumes more than a specified amount of RDBMS storage space across databases, tablespaces, groups of databases, and groups of tablespaces.

To create a databases-tablespaces quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager for Databases > Policy Management > Quotas**.

2. Select the level at which you want to apply the quota by right-clicking **Databases-Tablespaces**.
3. Select **Create Quota** from the pop-up menu. The Create Quota window opens.
4. In the **Description** field, type a brief description of the quota.
5. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP\_BY\_OWNER type profiles shown. Make sure that the selected profile is used in scans of all the databases and tablespaces chosen on the **Databases-Tablespaces** tab. This ensures that the quota is fairly enforced for all users.
6. Select users or user groups against which you want to apply the quota:
  - a. In the **Available** field, click the users or groups of users against which you want to apply the quota.
  - b. Click . The users or groups are moved to the **Current Selections** field.
7. Click the **Databases-Tablespaces** tab; the Databases-Tablespaces page opens.
8. Select the databases, tablespaces, database-tablespaces groups, computers, and computer groups on which you want to limit the storage usage of users and user groups.
  - a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups displayed in this field were created using **Data Manager for Databases > Monitoring > Groups**.
  - b. Click . The resources are moved to the **Current Selections** field.
9. Click the **When to Run** tab; the When to Run page opens.
10. In the **How often to run** area, schedule how often you want the quota to be checked. You can choose from the following actions:

#### **Run Now**

Check the quota immediately.

#### **Run Once at**

Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

**Note:** The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

#### **Run Repeatedly**

Check the quota repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
- Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark is displayed next to the days you select.

11. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
  - a. 9:00 AM of the time zone where the Data server resides
  - b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

**Use the timezone that the server runs in**

Select this option to use the time zone of the location where the Data server resides.

**Use this timezone**

Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field is activated.

12. Click the **Alert** tab; the Alert page opens.
13. In the **Triggering Condition** area, specify the event that will trigger the alert:
  - a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where *n* represents a value that you define in the **Value** field. No other conditions are available.
  - b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota is triggered.
  - c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.
14. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition. You can choose from the following actions:

**SNMP trap**

An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

**TEC / OMNIbus Event**

A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, and so on.

**Login notification**

The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

**Login ID**

The ID of the user who receives alert notification upon logging in to the system.

**Windows Event Log, UNIX Syslog**

The ID of the user who receives alert notification upon logging in to the system.

**Event Type (Windows Event Log only)**

The type of event that is recorded to the OS log.

**Facility (UNIX Syslog only)**

The location where the UNIX Syslog is sent. You can select User or Local.

**Level (UNIX Syslog only)**

The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

**Run script**

A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window are passed to the script when the alert is triggered



**E-mail** An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit e-mail** to change the actual email message.



15. Save the quota:
  - a. Click **File > Save**. The Save As window opens.
  - b. In the **Specify Quota name** field, type a name for the quota and click **OK**. The quota is displayed in the navigation tree pane

**Editing a storage usage quota**

Learn how to change the settings for a quota, including the users or user groups to which the quota applies, the resources on which the quota is set, and the amount set as the limit for storage usage.

To edit a quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager > Policy Management > Quotas**.
2. Expand the **User** or **OS User Group** node depending upon the type of quota you want to edit.
3. Expand the level of the quota you want to edit. For example, if you want to edit a computer quota, expand the **Computer** node.
4. Select the quota you want to edit. Information about the quota you want to edit appears in the right content pane.
5. In the **Description** field, type a brief description of the quota.
6. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP\_BY\_OWNER type profiles shown.
7. Select additional users or user groups against which you want to apply the quota:
  - a. In the **Available** field, click the users or groups of users against which you want to apply the quota. The groups displayed in this field were created using **Data Manager > Monitoring > Groups**
  - b. Click . The users or groups are moved to the **Current Selections** field.
8. Remove users or user groups previously included in the quota:
  - a. In the **Current Selections** field, click the users or groups of users that you want to remove from the quota.
  - b. Click . The users or groups are moved back to the **Available** field.

9. Click the **Computers** or **Filesystems** tab depending on quota you are editing; the Computers or Filesystems page opens.
10. Select additional filesystems, filesystem groups, computers, or computers groups on which you want to limit the storage usage of users and user groups.
  - a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups displayed in this field were created using **Data Manager > Monitoring > Groups**.
  - b. Click . The resources are moved to the **Current Selections** field.
11. Remove any filesystems, filesystem groups, computers, or computers groups that you do not want included in the quota.
  - a. In the **Current Selections** field, click the storage resources to remove.
  - b. Click . The resources are moved to the **Current Selections** field.
12. Click the **When to Run** tab; the When to Run page opens.
13. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

**Run Now**

Check the quota immediately.

**Run Once at**

Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

**Note:** The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

**Run Repeatedly**

Check the quota repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
- Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark is displayed next to the days you select.

14. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
  - a. 9:00 AM of the time zone where the Data server resides
  - b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

**Use the timezone that the server runs in**

Select this option to use the time zone of the location where the Data server resides.

### Use this timezone



Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field is activated.



15. Click the **Alert** tab; the Alert page opens.
16. In the **Triggering Condition** area, specify the event that will trigger the alert:
  - a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where *n* represents a value that you define in the **Value** field. No other conditions are available.
  - b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota is triggered.
  - c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.
17. Click **File > Save** to save your changes.

## Editing a storage usage quota for databases

Learn how to change the settings for a quota, including the users or user groups to which the quota applies, the resources on which the quota is set, and the amount set as the limit for storage usage.

To edit a quota, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager for Databases > Policy Management > Quotas**.
2. Expand the level of the quota you want to edit. For example, if you want to edit a instance quota, expand the **Instance** node.
3. Click the quota you want to edit. Information about the quota you want to edit appears in the right content pane.
4. In the **Description** field, type a brief description of the quota.
5. Select a profile from the **Profile** selection list to control which storage usage statistics are used when checking for quota violations. You can select any of the GROUP\_BY\_OWNER type profiles shown.
6. Select additional users or user groups against which you want to apply the quota:
  - a. In the **Available** field, click the users or groups of users against which you want to apply the quota. The groups displayed in this field were created using **Data Manager > Monitoring > Groups**
  - b. Click . The users or groups are moved to the **Current Selections** field.
7. Remove users or user groups previously included in the quota:
  - a. In the **Current Selections** field, click the users or groups of users that you want to remove from the quota.
  - b. Click . The users or groups are moved back to the **Available** field.
8. Click the **Databases-Tablespaces** or **Instances** tab depending on quota you are editing; the Databases-Tablespaces or Instances page opens.
9. Select additional databases, tablespaces, instances, computers, or computers groups on which you want to limit the storage usage of users and user groups.

- a. In the **Available** field, click the storage resources to which you want to apply the quota. The groups displayed in this field were created using **Data Manager for Databases > Monitoring > Groups**.
  - b. Click . The resources are moved to the **Current Selections** field.
10. Remove any databases, tablespaces, instances, computers, or computers groups that you do not want included in the quota.
    - a. In the **Current Selections** field, click the storage resources to remove.
    - b. Click . The resources are moved to the **Current Selections** field.
  11. Click the **When to Run** tab; the When to Run page opens.
  12. In the **How often to run** area, schedule how often you want the quota to be checked You can choose from the following actions:

#### **Run Now**

Check the quota immediately.

#### **Run Once at**

Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

**Note:** The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

#### **Run Repeatedly**

Check the quota repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
- Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark is displayed next to the days you select.

13. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
  - a. 9:00 AM of the time zone where the Data server resides
  - b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

#### **Use the timezone that the server runs in**

Select this option to use the time zone of the location where the Data server resides.

#### **Use this timezone**

Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field is activated.

14. Click the **Alert** tab; the Alert page opens.
15. In the **Triggering Condition** area, specify the event that will trigger the alert:

- a. In the **Condition** field, the default setting is **User Consumes More Than  $n$** , where  $n$  represents a value that you define in the **Value** field. No other conditions are available.
  - b. In the **Value** field, type a numeric value that represents the storage usage limit for the quota. If the storage entities selected for the quota (for example, users and user groups) exceed this limit, the quota is triggered.
  - c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.
16. Click **File** > **Save** to save your changes.

## Storage usage for Network Appliance quotas

Use Network Appliance quotas to import quotas from Network Appliance file servers and determine how and when you will be alerted to the hard limits defined in those quotas.

**Restriction:** You can only import quotas if the NetApp device is configured as a file server. It is not available if the NetApp device was added only as a CIMOM.

You can use the NAS support within IBM Tivoli Storage Productivity Center to work with quotas defined using the software provided with Network Appliance (NetApp) File Servers. Using the **Network Appliance** node under **Policy Management**, you can do the following:

- Use the **Schedules** node to define a Data Manager quota that will import NetApp quotas and alert you to when those quotas are close to being exceeded.
- Use the **Imported User Quotas** node to view the definitions of User quotas stored on the NetApp file servers against which you ran a Data Manager quota (defined in the **Schedules** node).
- Use the **Imported OS User Group Quotas** node to view the definitions of OS User Group quotas stored on the NetApp file servers against which you ran a Data Manager quota (defined in the **Schedules** node).

For example, you can define to be alerted when:

- A user or user group is close to reaching the hard limit defined within the corresponding NetApp Quota
- The space consumed on a QTree is approaching the hard limit defined within the corresponding NetApp Quota

Specifically, the Data Manager NAS component enables you to work with the following types of Network Appliance quotas:

- Quotas that limit the space that QTree directories can consume on a volume
- Quotas that limit the space consumed by specific users on a volume
- Quotas that limit the space consumed by specific user groups on a volume

### Quick Steps for Working with Network Appliance quotas

The following describes the general steps you must perform to work with Network Appliance Quotas:

1. Run a Data Manager scan against the Network Appliance file servers whose quotas you want to check. This ensures that Data Manager is using the latest statistics when determining quota violations.
2. Expand **Policy Management** > **Network Appliance Quotas**.




3. Right click the **Schedules** node and select **Create NetApp Quota Job** to define a quota job. The NetApp Quota window appears.
4. Use the tabs in the NetApp Quota window to specify:
  - the Network Appliance file servers whose quotas you want to import
  - when and how often to run the quota
  - the condition (% of hard limit) that will trigger an alert for the Network Appliance quotas and how you will be notified of that alert. For example, if you enter 90% for the alert condition, Data Manager will notify you whenever any of the quotas you are importing are within 10% of being violated.
5. Run the quota job from the **Schedules** node to import the Network Appliance quota definitions from the target NAS filer. When quotas are imported, the following occurs:
  - any User or OS User Group quotas violations are shown in the Quota Violation reports
  - any directory quota (Qtree) violations and definitions are displayed in **Alert Log >> Directory**.

### Creating a Network Appliance storage usage quota

Learn how to create a Network Appliance storage usage quota that imports Network Appliance quotas from filers and alert you to violations of those quotas.

Complete the following steps to:

- define the NetApp filers whose quotas you want to import
  - determine when and how often to run the quota that checks for user-defined violations (% of hard limit) of NetApp Quotas
  - specify the condition (% of hard limit) that will trigger an alert violation for a quota and how you will be notified of that alert
1. In the Navigation tree pane, expand **Data Manager > Policy Management > Network Appliance Quotas**.
  2. Right click the **Schedules** node.
  3. Select **New NetApp Quota Job** from the pop-up menu. The content pane displays the What to PROBE page.
  4. In the **Description** field, type a brief description of the quota.
  5. Select the NetApp filers against which you want to run the quota job:
    - a. In the **Available** field, click the NetApp filers or NetApp filer groups against which you want to apply the quota.
    - b. Click . The NetApp filers or NetApp filer groups are moved to the **Current Selections** field.
  6. Click the **When to Run** tab; the When to Run page opens.
  7. In the **How often to run** area, schedule how often you want the quota to be checked. You can choose from the following actions:

#### Run Now

Check the quota immediately.

#### Run Once at

Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

**Note:** The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

### Run Repeatedly

Check the quota repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
  - Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
  - Indicate which days on which run the job. A check mark is displayed next to the days you select.
8. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
- a. 9:00 AM of the time zone where the Data server resides
  - b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

### Use the timezone that the server runs in

Select this option to use the time zone of the location where the Data server resides.

### Use this timezone

Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field is activated.

9. Click the **Alert** tab; the Alert page opens.
10. In the **Triggering Condition** area, specify the event that will trigger the alert:
- a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where *n* represents a value that you define in the **Value** field. No other conditions are available.
  - b. In the **Value** field, enter the value that represents the percentage of a NetApp's hard limit at which the alert is triggered. The hard limit represents the most amount of space a user or directory (Qtree) can consume on a NAS filer. For example, if the hard limit defined within a NetApp Quota for a user group is 150 MB, and you enter **90** in the **Value** field, the Data Manager quota will report a violation when that user group consumes 135 or more megabytes.
  - c. In the **Value Units** field, the only value units available is the % of Hard Limit.
11. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition. You can choose from the following actions:

### SNMP trap

An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

### TEC / OMNIBus Event

A IBM Tivoli Enterprise Console event is sent when the condition

occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, and so on.

**Login notification**

The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

**Login ID**

The ID of the user who receives alert notification upon logging in to the system.

**Windows Event Log, UNIX Syslog**

The ID of the user who receives alert notification upon logging in to the system.

**Event Type (Windows Event Log only)**

The type of event that is recorded to the OS log.

**Facility (UNIX Syslog only)**

The location where the UNIX Syslog is sent. You can select User or Local.

**Level (UNIX Syslog only)**

The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

**Run script**

A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window are passed to the script when the alert is triggered

**E-mail** An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit e-mail** to change the actual email message.

12. Save the quota:
  - a. Click **File > Save**. The Save As window opens.
  - b. In the **Specify NetApp Quota job name** field, type a name for the quota and click **OK**. The quota is displayed in the Navigation Tree pane

**Viewing user quotas imported from Network Appliance filers**

Learn how to view the definitions of user quotas stored on the NetApp filers against which you ran a Data Manager quota.

Before you can import Network Appliance user quotas and view them through **Data Manager > Policy Management > Network Appliance Quotas > Imported User Quotas**, you must:

1. Run a scan against the Network Appliance filers whose user quotas you want to import.
2. Use **Policy Management > Network Appliance Quotas > Schedules** to define and run a quota checking job against the NetApp filers whose user quotas you want to import.

Complete the following steps to view imported Network Appliance user quotas:

1. In the Navigation tree pane, expand **Data Manager > Policy Management > Network Appliance Quotas > Imported User Quotas**.
2. Click the NetApp Quota whose definition you want to view. The content pane displays information about the NetApp user quota.
3. In the Users page of the Edit *quota\_name* window, view the following information about imported Network Appliance user quotas:

**Hard Limit**

The hard limit for used space defined within the NetApp quota.

**Current Selections**

The users against which the NetApp quota applies.

4. In the Filesystems page of the Edit *quota\_name* window, view the filesystems to which the NetApp Quota applies.

### Viewing OS user quotas imported from Network Appliance filers

Learn how to view the definitions of OS User Quotas stored on the NetApp filers against which you ran a Data Manager quota.

Before you can import Network Appliance OS user quotas and view them through **Data Manager > Policy Management > Network Appliance Quotas > Imported OS User Group Quotas**, you must:

1. Run a scan against the Network Appliance filers whose OS user quotas you want to import.
2. Use **Policy Management > Network Appliance Quotas > Schedules** to define and run a quota checking job against the NetApp filers whose OS user quotas you want to import.

Complete the following steps to view imported Network Appliance OS user quotas:

1. In the Navigation tree pane, expand **Data Manager > Policy Management > Network Appliance Quotas > Imported OS User Quotas**.
2. Click the NetApp Quota whose definition you want to view. The content pane displays information about the NetApp OS user quota.
3. In the OS Users page of the Edit *quota\_name* window, view the following information about imported Network Appliance OS user quotas:

**Hard Limit**

The hard limit for used space defined within the NetApp quota.

**Current Selections**

The OS users against which the NetApp quota applies.

4. In the Filesystems page of the Edit *quota\_name* window, view the filesystems to which the NetApp Quota applies.

### Viewing directory (QTree) quotas and violations

Learn how to view the definitions of user quotas stored on the Network Appliance filers against which you ran a Data Manager quota.

Before you can view QTree quotas and violations, perform the following steps:

1. Use **Policy Management > Network Appliance Quotas > Schedules** to define and run a quota checking job against the Network Appliance filers containing the QTree quotas you want to import.
2. Create and run a scan job that includes the following objects:
  - the NetApp filers containing the QTrees whose quotas you want to import

- the directory group named TPCUser.NetApp QTree Group
3. Use **Policy Management > Network Appliance Quotas > Schedules** to rerun the quota checking job you defined in step 1.

Complete the following steps to view a QTree quota violation and definition:

1. In the Navigation tree pane, expand **Alerting > Alert Log**.
2. Click the **Directory** node. The Alert History — All Directories window appears. This window lists all the alerts triggered by directory and QTree quota violations.
3. To view details about a violation you can do either of the following:
  - Double-click the entry representing the QTree violation to view details about that violation and the QTree quota definition.
  - Right click the entry representing the QTree violation and select **Edit Definition** from the pop-up menu to view the QTree quota definition.

**Note:** You cannot edit a QTree quota definition using Data Manager. To edit a QTree quota definition you must use the software provided with your Network Appliance filer.


## Setting file constraints

Use constraints to define the acceptable and unacceptable file types, file sizes, and file owners for a NAS filer or a set of NAS filers in your environment. You can also use constraints to request an IBM Tivoli Storage Manager archive and backup of the largest violating files identified by a constraint.

### Creating a file constraint

Learn how to create a constraint that defines the acceptable and unacceptable file types, file sizes, and file owners for a computer or a set of computers in your environment.

To create a constraint, complete the following steps:



1. In the Navigation tree pane, expand **Data Manager > Policy Management**.
2. Right-click **Constraints**, and click **Create Constraint**. The Create Constraint window opens.
3. In the **Description** field, type a brief description of the job.
4. Select file systems, file system groups, computers, and computer groups against which to apply the constraint:
  - a. In the **Available** field, click the storage resources against which you want to apply the constraint. Select individual filesystems and computers, as well as file system and computer groups that you have previously defined using the group facility. Note that cluster resource groups in a PowerHA or MSCS clustered environment are listed as computers.
  - b. Click . The storage resources are moved to the **Current Selections** field.
5. Click the **File Types** tab; the File Types page opens. Use this page to select the file types that you want to allow or forbid.
6. Click one of the following:

#### Forbid file matching these patterns

Do *not* allow files of the specified type on the storage resources selected on the Filesystems page.

### Allow ONLY file matching these patterns

Allow files *only* of the specified type on the storage resources selected on the Filesystems page.


7. Click the file types that you would like to allow or forbid in the **Choose a file type** list box. The file types listed in this box are predefined in the system.
  - a. In the **Choose a file type:** field, click the file types that you would like to allow or forbid.
  - b. Click . The file types are moved to the **Forbidden files** field.
8. Optionally, create a custom file type:
  - a. Enter the pattern in the **Or enter a pattern** field.
  - b. Click . The file type is moved to the **Forbidden files** field.
9. Click the **Users** tab; the Users page opens. Use this page to allow or forbid files owned by specific users
10. Click one of the following:

#### Forbid file owned by selected users

Do *not* allow files owned by the specified users.

#### Allow ONLY file owned by the specified owner

Allow files owned by the specified users only.

11. Select the users whose files you want to allow or forbid:
  - a. In the **Available Users** field, click the users you want to include.
  - b. Click . The users are moved to the **Forbidden Users** field.
12. Click the **Options** tab; the Options page opens. Use this page to define alerts for files based on characteristics other than file type and file owner, and to view a where clause that will be used as a filter for locating files that violate the constraint definition. This page contains the following information:

#### Violating File Limits

Indicate the maximum number of violating filenames that should be kept per agent.

#### number of days since last access greater than

Click this box and enter a length of time (in days) in which files have not been accessed. When a constraint job locates a file that has not been accessed in the length of time you indicated in this field, an alert is triggered.

#### bigger than ##

Check this box and enter a file size that should trigger an alert. When a constraint job locates a file of the size (or greater) that you indicated in this field, an alert is triggered. Select the unit of measurement for the value you entered in the list box to the right. The valid choices are: bytes, **KB** (kilobytes), **MB** (megabytes), **GB** (gigabytes), and **TB** (terabytes).

#### set-UID root

Check this box trigger alerts on files that have the SET-UID root bit on (because this can be a security consideration).

#### not backed up (Windows only) and number of days since last modification greater than

Click this box and enter the number of days since a file was last modified and not backed up. When the constraint job locates a file

that was modified and not backed up in the amount of days that you indicate in this field, an alert is triggered.

#### **File Filter text**

This section displays a *where* clause that will be used as a filter for locating files that violate the Constraint definition. The choices you make on this tab and the choices you made on the Users and File Types pages are reflected in the content of this *where* clause. To edit the *where* clause filter text, click **Edit Filter**. This will allow you to create new where conditions and groupings using the where clause interface.

13. Click the **Alert** tab; the Alert page opens. Use the Alert tab to define the alerting criteria for a constraint.
14. In the **Triggering Condition** area, specify the event that will trigger the constraint alert:
  - a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where *n* represents a value that you define in the **Value** field. No other conditions are available.
  - b. In the **Value** field, type a numeric value that represents the amount of space a file must consume to violate the constraint.
  - c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.

Once Data Manager has found enough files that violate the condition (i.e., that exceed the total size amount you specified in this section), the constraint is considered violated and the alert is triggered. When the alert is triggered, the actions that you specify in the **Triggered Actions** section are taken.

15. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition. You can choose from the following actions:

#### **SNMP trap**

An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

#### **TEC / OMNIbus Event**

A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, and so on.

#### **Login notification**

The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

#### **Login ID**

The ID of the user who receives alert notification upon logging in to the system.

#### **Windows Event Log, UNIX Syslog**

The ID of the user who receives alert notification upon logging in to the system.

#### **Event Type (Windows Event Log only)**

The type of event that is recorded to the OS log.

**Facility (UNIX Syslog only)**

The location where the UNIX Syslog is sent. You can select User or Local.

**Level (UNIX Syslog only)**

The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

**Run script**

A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window are passed to the script when the alert is triggered

**E-mail** An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit e-mail** to change the actual email message.



## 16. Save the constraint:

- a. Click **File > Save**. The Save As window opens.
- b. In the **Specify Constraint name** field, type a name for the constraint and click **OK**. The constraint is displayed in the navigation tree pane. The constraint is automatically checked whenever an appropriate scan job is run.

**Editing a constraint**

Edit a constraint to: change the acceptable and unacceptable file types, file sizes, and file owners, add or remove users and user groups to which the constraint applies, edit the constraint options, and change the alert settings.

To edit a constraint, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager > Policy Management > Constraints**.
2. Click the constraint you want to edit. Information about the constraint you want to edit appears in the right content pane.
3. Add or remove the filesystems, filesystem groups, computers, and computer groups against which to apply the constraint:
  - a. In the **Available** field, click the storage resources that you want add. Click . The users or groups are moved to the **Current Selections** field.
  - b. In the **Current Selections** field, click the storage resources that you want remove from the constraint. Click . The users or groups are moved to the **Available** field.
4. Click the **File Types** tab; the File Types page opens. Use this page to edit the file types that you want to allow or forbid.
5. Click one of the following:

**Forbid file matching these patterns**



Do *not* allow files of the specified type on the storage resources selected on the Filesystems page.

**Allow ONLY file matching these patterns**

Allow files *only* of the specified type on the storage resources selected on the Filesystems page.

6. Click the file types that you would like to allow or forbid in the **Choose a file type** list box. The file types listed in this box are predefined in the system.




- a. In the **Choose a file type:** field, click the file types that you would like to allow or forbid.
  - b. Click . The file types are moved to the **Forbidden files** field.
7. Optionally, create a custom file type:
    - a. Enter the pattern in the **Or enter a pattern** field.
    - b. Click . The file type is moved to the **Forbidden files** field.
  8. Click the **Users** tab; the Users page opens. Use this page to allow or forbid files owned by specific users
  9. Click one of the following:

**Forbid file owned by selected users**

Do *not* allow files owned by the specified users.

**Allow ONLY file owned by the specified owner**

Allow files owned by the specified users only.

10. Select the users whose files you want to allow or forbid:
  - a. In the **Available Users** field, click the users you want to include.
  - b. Click . The users are moved to the **Forbidden Users** field.
11. Click the **Options** tab; the Options page opens. Use this page to edit alerts for files based on characteristics other than file type and file owner, and to view a where clause that will be used as a filter for locating files that violate the constraint definition. This page contains the following information:

**Violating File Limits**

Indicate the maximum number of violating filenames that should be kept per agent.

**number of days since last access greater than**

Click this box and enter a length of time (in days) in which files have not been accessed. When a constraint job locates a file that has not been accessed in the length of time you indicated in this field, an alert is triggered.

**bigger than ##**

Check this box and enter a file size that should trigger an alert. When a constraint job locates a file of the size (or greater) that you indicated in this field, an alert is triggered. Select the unit of measurement for the value you entered in the list box to the right. The valid choices are: bytes, **KB** (kilobytes), **MB** (megabytes), **GB** (gigabytes), and **TB** (terabytes).

**set-UID root**

Check this box trigger alerts on files that have the SET-UID root bit on (because this can be a security consideration).

**not backed up (Windows only) and number of days since last modification greater than**

Click this box and enter the number of days since a file was last modified and not backed up. When the constraint job locates a file that was modified and not backed up in the amount of days that you indicate in this field, an alert is triggered.

**File Filter text**

This section displays a *where* clause that will be used as a filter for locating files that violate the Constraint definition. The choices you make on this tab and the choices you made on the Users and File

Types pages are reflected in the content of this *where* clause. To edit the *where* clause filter text, click **Edit Filter**. This will allow you to create new where conditions and groupings using the where clause interface.

12. Click the **Alert** tab; the Alert page opens. Use the Alert tab to define the alerting criteria for a constraint.
13. In the **Triggering Condition** area, specify the event that will trigger the constraint alert:
  - a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where *n* represents a value that you define in the **Value** field. No other conditions are available.
  - b. In the **Value** field, type a numeric value that represents the amount of space a file must consume to violate the constraint.
  - c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.

Once Data Manager has found enough files that violate the condition (i.e., that exceed the total size amount you specified in this section), the constraint is considered violated and the alert is triggered. When the alert is triggered, the actions that you specify in the **Triggered Actions** section are taken.

14. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition. You can choose from the following actions:

**SNMP trap**

An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

**TEC / OMNIbus Event**

A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, and so on.

**Login notification**

The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

**Login ID**

The ID of the user who receives alert notification upon logging in to the system.

**Windows Event Log, UNIX Syslog**

The ID of the user who receives alert notification upon logging in to the system.

**Event Type (Windows Event Log only)**

The type of event that is recorded to the OS log.

**Facility (UNIX Syslog only)**

The location where the UNIX Syslog is sent. You can select User or Local.

**Level (UNIX Syslog only)**

The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

**Run script**

A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window are passed to the script when the alert is triggered

**E-mail** An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit e-mail** to change the actual email message.

15. Click **File > Save** to save your changes.

**Requesting an archive or backup**

Request an IBM Tivoli Storage Manager archive and backup of the largest violating files identified by a constraint. Tivoli Storage Manager protects your data from hardware failures and other errors by storing backup and archive copies of data on auxiliary storage.

Data Manager provides a method within the alerting function of a constraint where you can automatically start a Tivoli Storage Manager archive or backup to run against the files that violate that constraint. For example, you can do the following actions:

- Define a constraint to identify the .vbs files within your organization that are larger than 2 megabytes. Other file types include .avi, .dll, .doc, .exe, .gif, .jpg, .mp3, .rtf, .txt., wsh, and .zip. You can also define a pattern to use when identifying files that need to be archived or backed up.
- Define a constraint that runs Tivoli Storage Manager to archive and then delete those files

**Before you begin**

Before you can use the Tivoli Storage Manager archive and backup functions with Data Manager constraints, you must ensure that the following conditions are met:


- Your user ID has Administrator authority within Tivoli Storage Productivity Center.
- Tivoli Storage Manager client software is installed on each of the hosts where a Data agent or Storage Resource agent is used to perform archive and backup functions.
- (For SAN attached storage only) Tivoli Storage Manager client software must be installed on a host computer where a Data agent or Storage Resource agent is configured to act as the scanning agent for NAS storage.
- A Tivoli Storage Manager client must be configured to run without a password. This condition is typically done by either including the password in the client options (dsm.opt) file, or using the options `PASSWORDACCESS = GENERATE` in the client options file.
- The Tivoli Storage Manager client must be version 5.1 or higher

**Defining an archive or backup in a constraint:**



Learn how to request an IBM Tivoli Storage Manager archive and backup of the largest violating files identified by a constraint. Request a Tivoli Storage Manager archive and backup on the alert page of a constraint definition.

**Create a new constraint:**


1. In the Navigation tree pane, expand **Data Manager > Policy Management**.

2. Right-click **Constraints**, and click **Create Constraint**. The Create Constraint window opens.
3. Select filesystems, filesystem groups, computers, and computer groups against which to apply the constraint:
  - a. In the **Available** field, click the storage resources against which you want to apply the constraint. Select individual filesystems and computers, as well as file system and computer groups that you have previously defined using the group facility. Note that cluster resource groups in a PowerHA or MSCS clustered environment are listed as computers.
  - b. Click . The users or groups are moved to the **Current Selections** field.
4. Click the **File Types** tab; the File Types page opens. Use this page to select the file types that you want to allow or forbid.
5. Click one of the following:
 

**Forbid file matching these patterns**  
Do *not* allow files of the specified type on the storage resources selected on the Filesystems page.

**Allow ONLY file matching these patterns**  
Allow files *only* of the specified type on the storage resources selected on the Filesystems page.
6. Click the file types that you would like to allow or forbid in the **Choose a file type** list box. The file types listed in this box are predefined in the system.
  - a. In the **Choose a file type:** field, click the file types that you would like to allow or forbid.
  - b. Click . The file types are moved to the **Forbidden files** field.
7. Optionally, create a custom file type:
  - a. Enter the pattern in the **Or enter a pattern** field.
  - b. Click . The file type is moved to the **Forbidden files** field.
8. Click the **Users** tab; the Users page opens. Use this page to allow or forbid files owned by specific users
9. Click one of the following:
 

**Forbid file owned by selected users**  
Do *not* allow files owned by the specified users.

**Allow ONLY file owned by the specified owner**  
Allow files owned by the specified users only.
10. Select the users whose files you want to allow or forbid:
  - a. In the **Available Users** field, click the users you want to include.
  - b. Click . The users are moved to the **Forbidden Users** field.
11. Click the **Options** tab; the Options page opens. Use this page to define alerts for files based on characteristics other than file type and file owner, and to view a where clause that will be used as a filter for locating files that violate the constraint definition. This page contains the following information:
 

**Violating File Limits**  
Indicate the maximum number of violating filenames that should be kept per agent.

**number of days since last access greater than**

Click this box and enter a length of time (in days) in which files have not been accessed. When a constraint job locates a file that has not been accessed in the length of time you indicated in this field, an alert is triggered.

**bigger than ##**

Check this box and enter a file size that should trigger an alert. When a constraint job locates a file of the size (or greater) that you indicated in this field, an alert is triggered. Select the unit of measurement for the value you entered in the list box to the right. The valid choices are: bytes, **KB** (kilobytes), **MB** (megabytes), **GB** (gigabytes), and **TB** (terabytes).

**set-UID root**

Check this box trigger alerts on files that have the SET-UID root bit on (because this can be a security consideration).

**not backed up (Windows only) and number of days since last modification greater than**

Click this box and enter the number of days since a file was last modified and not backed up. When the constraint job locates a file that was modified and not backed up in the amount of days that you indicate in this field, an alert is triggered.

**File Filter text**

This section displays a *where* clause that will be used as a filter for locating files that violate the Constraint definition. The choices you make on this tab and the choices you made on the Users and File Types pages are reflected in the content of this *where* clause. To edit the *where* clause filter text, click **Edit Filter**. This will allow you to create new where conditions and groupings using the where clause interface.

12. Click the **Alert** tab; the Alert page opens. Use the Alert tab to define the alerting criteria for a constraint.
13. In the **Triggering Condition** area, specify the event that will trigger the constraint alert:
  - a. In the **Condition** field, the default setting is **User Consumes More Than *n***, where *n* represents a value that you define in the **Value** field. No other conditions are available.
  - b. In the **Value** field, type a numeric value that represents the amount of space a file must consume to violate the constraint.
  - c. In the **Value Units** field, select the unit of measurement: kilobytes, megabytes, or gigabytes.

Once Data Manager has found enough files that violate the condition (i.e., that exceed the total size amount you specified in this section), the constraint is considered violated and the alert is triggered. When the alert is triggered, the actions that you specify in the **Triggered Actions** section are taken.

14. In the **Triggered Actions** area, select **Archive/Backup** and click **Define**. The Archive / Backup Options window appears. Click one of the following:

**Archive**

Perform an archive on the files that violate a constraint. Archives are copies of data from online storage copies to offline storage. If you select the Archive option, you can also check the **Delete After**

**Successful Archival** box to automatically delete the files that violated the constraint after the files have been successfully archived.

### **Selective Backup**

Perform a selective backup of the files that violate a constraint. A selective backup enables you to back up specified files. These files are not excluded in the include-exclude list and meet the requirement for serialization in the backup copy group of the management class assigned to each file.

### **Incremental Backup**

Perform an incremental backup of the files that violate a constraint. An incremental backup enables you to back up new or changed files or directories from a client domain or from specified directories or files. These directories or files are not excluded in the include-exclude list and meet the requirements for frequency, mode, and serialization as defined by a backup copy group of the management class assigned to each file.

### **Tivoli Storage Manager Command Options**

Enter additional that will be passed on to the Tivoli Storage Manager `dsmc` command. You can enter up to 512 bytes of text.

15. Save the constraint:
  - a. Click **File > Save**. The Save As window opens.
  - b. In the **Specify Constraint name** field, type a name for the constraint and click **OK**. The constraint is displayed in the navigation tree pane. The constraint is automatically checked whenever an appropriate scan job is run.

See the *IBM Tivoli Storage Manager Client User's Guide* for more information about each of the TSM actions you can select when defining triggered actions for a constraint.

## **Using filesystem extension**

Use filesystem extensions to create additional space in the local filesystems of managed hosts. You can extend filesystems manually, or set up policies to do it automatically.

Use **Data Manager > Policy Management > Filesystem Extension** to create additional space in the local filesystems of managed hosts either manually or automatically through a policy. Policies can be configured to extend filesystems at a specified time, or when utilization reaches a specified threshold. For managed hosts that have access to an TotalStorage Enterprise Storage Server, you can also allocate additional LUNs when there is not enough space to extend file systems in the local volume group.

Filesystem extension is supported for JFS file systems running on AIX 5.1 and VxFS file systems running on Solaris 2.8. The AIX JFS2 file system is not supported.

**Note:** Filesystem extension/LUN provisioning jobs are not supported under Solaris 9.

## How Filesystem Extension works

Filesystem extension is defined as a policy. The policy specifies the following:

- Which filesystems will be eligible for extension
- How the filesystems will be extended
- Whether LUN provisioning will be enabled, and how it will work
- When to enforce the policy
- How alerting will be handled

You can also choose to have no extension performed when the policy is applied. Instead, any policy actions that would have been performed are written to a log file. This feature, called LOG ONLY mode, can be used to preview the results of a policy before extending any file systems.

## LUN Provisioning

By default, IBM Tivoli Storage Productivity Center will attempt to extend a filesystem in its local volume group. If there is not enough space in the local volume group and an TotalStorage Enterprise Storage Server subsystem is available, additional LUNs can be provisioned for filesystem extension. Provisioning consists of creating and formatting a new LUN according to the parameters specified in the filesystem extension policy, and assigning the new LUN to the managed host. You can enable and configure LUN provisioning using the Provisioning tab in the Filesystem Extension window.

The Data server uses the TotalStorage Enterprise Storage Server Common Information Model/Object Manager (CIMOM) to interact with TotalStorage Enterprise Storage Server subsystems. The CIMOM is installed as an agent on a host server or workstation in your network. In many cases, CIMOMs can be automatically discovered, using Service Location Protocol (SLP). The ability to discover a CIMOM depends on its location and the operating system on which it is installed.

When a CIMOM is discovered or manually added, you must provide access information for it, including the username, password, and location of the TrustStore certificate file. The certificate file is used to ensure secure communication between the Data Manager server and the CIMOM. The certificate file is created when you install and set up the CIMOM. The certificate file for each CIMOM must be copied to the Data Manager server machine. For information about supported versions of the CIMOM, see the Tivoli Storage Productivity Center support website at [http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli\\_Storage\\_Productivity\\_Center\\_Standard\\_Edition](http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Productivity_Center_Standard_Edition). For more information and detailed instructions for setting up CIMOM access, see the CIMOM Logins section in the *IBM Tivoli Storage Productivity Center for Data Installation and Configuration Guide*.

**Note:** To use a CIMOM with Data Manager, eFix 1.1.0.2 must be installed on the CIMOM machine. This patch is available from the Tivoli Storage Productivity Center Technical Support Website.

## Creating a file system extension

Learn how to create a file system extensions policy that automatically creates additional space in the local file systems of managed hosts.

To create a file system extension policy, complete the following steps:


1. In the Navigation tree pane, expand **Data Manager > Policy Management**.
2. Right-click **Filesystem Extension**, and click **Create Filesystem Extension Rules**. The Create Filesystem Extension Rules window opens and the Filesystems page is displayed. Use this page to select the file systems to which this policy will apply. Filesystem extension does not currently support the use of predefined file system groups. Only individual file systems can be selected when defining extension policy. However, multiple file systems can be specified for the same policy.

A file system can only have one policy assigned to it. If a file system is already associated with a policy, the policy name will be displayed next to the file system name. If you select this file system for the extension policy, the previous policy assignment will be overridden when you save the new policy.

**Note:** When a new file system is added, it will appear in the Available list box only after a Probe job is run.

The file systems tree view includes a file system extension preview function. File systems that would be extended under the currently defined policy are displayed with a green arrow icon. If a file system would qualify for extension but has reached the maximum size specified by the policy, it is displayed with a red arrow icon. Click the Refresh button to update the tree view.

**Note:** File system extension supports the use of directly attached Serial Storage Architecture (SSA) devices, such as the IBM Tivoli Storage 7133 Serial Disk System. This support is provided for AIX versions 5.1, 5L, and 5.2. For operating systems that connect to SSA devices by emulating SCSI, Data Manager will recognize the SSA device as a SCSI drive.

3. In the **Description** field, type a brief description of the job.
4. Select the file systems for the policy:
  - a. In the **Available** field, click the file systems you want to include.
  - b. Click . The file systems are moved to the **Current Selections** field.
5. Click the **Extension** tab; the Extension page opens. Use this page to specify how file systems will be extended. You can choose to extend file systems:
  - By a specific amount of space
  - By a percentage of file system capacity
  - By increasing free space to a specific size
  - By increasing free space to a percentage of file system capacity
  - By increasing capacity to a specific size

Select the option you want to use for this policy. You can also set a maximum capacity size if necessary. If you specify a maximum size, a file system extension job that would increase a file system beyond that size will fail.

**Note:** The actual amount of file system extension can vary, depending on your environment. For example, if 5GB physical partitions are set up for a volume group, its file systems will be extended in 5GB increments. If you specify a smaller amount, it will be rounded up to the physical partition size.

As another example, software striping and software RAID will claim part of the space created when file systems are extended. In this case, file system extension will appear to create slightly less space than specified in the policy.

If you are using LUN provisioning, you can specify the size of new LUNs by using the Create LUNs that are at least option on the Provisioning tab and



one of the Amount to Extend options on the Extension tab. If the Amount to Extend value is greater than or equal to the preferred minimum LUN size, 1% more space will be added to the LUN to cover overhead requirements. If you choose to extend file systems regardless of remaining free space, the policy options you specify will be applied to all selected file systems when the policy is enforced. To use policy to automatically manage file system capacity, you must do the following:

- Use the **When to Run** page to set the policy to run after every probe or scan
- Use the **Extension** page to specify an amount of free space to be used as a condition for extension. You can do this two ways:
  - Select **Make Freespace** and specify an amount. Then select **Extend filesystems regardless of remaining free space**. When free space falls below the amount you specify, the difference in space will be added. This option can be used to maintain file system free space at a specific level.
  - Select either **Add** or **Make Freespace** and specify an amount. Then select **Extend filesystems when free space is less than** and specify an amount. The amount you specify will be used to trigger the **Amount to Extend** action you selected. This option can be used to add extra space to a file system if free space reaches a specified level. For example, if free space reaches 10% or less, it could be increased to 20%, or 5GB could be added to the total file system capacity.

**Note:** The following restrictions apply when you create or edit a file system extension policy:

- If you select both **Make Freespace** and **Extend filesystems when freespace is less than**, you must choose the same type of extension for each. If you specify a percentage amount for one, you must specify a percentage amount for the other. If you specify an amount expressed in megabytes for one, you must specify an amount in either megabytes or gigabytes for the other.
- If you select both **Make Freespace** and **Extend filesystems when freespace is less than**, the extension amount you specify for **Make Freespace** must be larger than the amount you specify for **Extend filesystems when freespace is less than**.
- If you select **Make Capacity**, you cannot select **Extend filesystems when freespace is less than**. Because the **Make Capacity** option extends capacity regardless of actual utilization, it is not intended to be used for the automated maintenance of file system capacity. You can use this option to extend file system capacity when necessary by selecting **Run Now** or **Run Once** at on the **When to Run** page.

You can also choose to have no file systems extended when the policy is applied. Instead, any policy actions that would have been performed are written to the log file. This feature, called LOG ONLY mode, can be used to preview the results of a policy over time before actually extending any file systems.

6. Click the **Provisioning** tab; the Provisioning page opens. Use the Provisioning tab to enable LUN provisioning and specify how it will be performed. Provisioning is an optional feature that is not required for file system extension. If you enable LUN provisioning, two things will happen when there is not enough space in a volume group to extend a file system:
  - A new LUN is created to add space to the volume group, according to the policy specified on the **Provisioning** tab.

- The file system is extended within the volume group, according to the policy specified on the **Extension** tab.

For example, a volume group might first be extended by 100Mb, and then one or more file systems within it might be extended by 10MB each.

Data Manager relies on information stored by a CIM/OM to interact with TotalStorage Enterprise Storage Server subsystems. To implement provisioning, CIM/OM access information must be defined to Data Manager. See the CIM/OM Logins section in the Configuration and Getting Started Guide for more information. After CIM/OM access is configured, a Discovery job must be run to identify TotalStorage Enterprise Storage Server subsystems.

**Note:** LUNs can be provisioned for file system hosts running Solaris, but the hosts must be specially configured to avoid a reboot after provisioning. See the appendix on configuring LUN provisioning for Solaris in the IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide for instructions.

If you choose to enable automatic LUN provisioning, you must choose a model and location for new LUNs. Data Manager models new LUNs on existing LUNs. Modeled characteristics include CIM/OM access parameters, as well as connectivity and RAID type. You must specify how far Data Manager will look for a model LUNs, as well as limit where it will create new LUNs on the TotalStorage Enterprise Storage Server.

**Note:** To perform LUN provisioning, at least one TotalStorage Enterprise Storage Server LUN must be currently assigned to the volume group of the file system you want to extend, or to the managed host associated with the file system you want to extend. Select one of the following options to limit the scope of LUN modeling:

- Model new LUNs on others in the volume group of the file system being extended. If you select this option, Data Manager will only search for a model within existing LUNs in the volume group of the file system being extended.
- Model new LUNs on others on the same host as the file system being extended. If you select this option, Data Manager will first search for a model within existing LUNs in the volume group of the file system being extended, before looking further on the same host.

New LUNs can be created in the same Storage Pool as the file system being extended, or in any Storage Pool in the TotalStorage Enterprise Storage Server of the file system being extended. Select one of the following options to limit where new LUNs will be created:

- Provision new LUNs within the same Storage Pool as the file system being extended. For TotalStorage Enterprise Storage Server, a Storage Pool is a Volume Space. If you are using RAID, a Volume Space is equal to one RAID rank, which consists of a Disk Group (8 physical disks).
- Provision new LUNs within the same Storage System as the file system being extended. If you select this option, new LUNs can be created in any Storage Pool in the TotalStorage Enterprise Storage Server. Data Manager will first attempt to provision a new LUN in the same Storage Pool as the file system being extended, before looking further in the same Storage System.

You can optionally specify a preferred minimum size for new LUNs by selecting the **Create LUNs that are at least** option. If you select this option, LUNs will not be created any smaller than the size you specify. If you do not

specify a preferred minimum size, the Amount to Extend specified for the policy will be used. For TotalStorage Enterprise Storage Server subsystems, the default minimum LUN size increment is 100 MB. New LUNs will be rounded up to satisfy this requirement. If the Amount to Extend value is greater than or equal to the preferred minimum LUN size you specify, 1% more space will be added to the LUN to cover overhead requirements. See Table 23.

**Note:** Data Manager will only provision new LUNs. Existing unassigned LUNs will not be reused. Similarly, if provisioning fails after LUNs have been created, the new LUNs will not be deleted or reused. This limitation protects against accidentally overwriting existing data that might be associated with currently unassigned LUNs.

Table 23 shows the possible outcomes when a new LUN is provisioned.

Table 23. LUN Provisioning Results

| Description                                                                                                            | Amount to Extend | TotalStorage Enterprise Storage Server minimum for new LUNs | Create LUNs that are at least (preferred minimum LUN size) | Actual LUN size that will be created | Actual amount by which the file system will be extended |
|------------------------------------------------------------------------------------------------------------------------|------------------|-------------------------------------------------------------|------------------------------------------------------------|--------------------------------------|---------------------------------------------------------|
| Minimum LUN size exceeds amount to extend                                                                              | 20Mb             | 100 Mb                                                      | 200Mb                                                      | 200Mb                                | 20Mb                                                    |
| Amount to extend equals preferred minimum LUN size                                                                     | 200Mb            | 100 Mb                                                      | 200Mb                                                      | 202Mb (1% added for overhead)        | 200Mb                                                   |
| Minimum LUN size exceeds amount to extend                                                                              | 100Mb            | 100 Mb                                                      | 200Mb                                                      | 200Mb                                | 100Mb                                                   |
| Amount to extend exceeds preferred minimum LUN size, but both are below TotalStorage Enterprise Storage Server minimum | 75Mb             | 100 Mb                                                      | 50Mb                                                       | 100Mb                                | 75Mb                                                    |
| Minimum LUN size not specified                                                                                         | 200Mb            | 100 Mb                                                      | Not specified                                              | 202Mb (1% added for overhead)        | 200Mb                                                   |

**Note:** IBM TotalStorage Enterprise Storage Server and theTotalStorage Enterprise Storage Server CIM/OM use multiples of 1000 bytes to represent kilobytes, megabytes, and gigabytes. Data Manager uses multiples of 1,024 bytes. Data Manager automatically converts this difference when it discovers TotalStorage Enterprise Storage Server capacity and utilization. If you are using IBM Tivoli Storage TotalStorage Enterprise Storage Server Expert or TotalStorage Enterprise Storage Server Specialist, the storage information they show will differ from what is shown by Data Manager.

Due to AIX Volume Manager limitations on AIX Striped Volumes, provisioning of new LUNs cannot be done. If the policy definition for a striped file system indicates to Enable Automatic LUN Provisioning on the Provisioning tab, then any LUNs that are created will be added to the Volume Group but will not be usable for Extension. If provisioning is attempted on an AIX Striped Volume, then you will see the message: STA0301W: To prevent

runaway condition, we have removed the file system file system\_name from the policy, please fix the problem and add the file system back to the policy.

To avoid this problem, make sure the **Enable Automatic LUN Provisioning** option is disabled for Filesystem Extension policy definitions that have striped file systems on AIX.

7. Click the **When to Run** tab; the When to Run page opens. Use this page to specify when to apply the file system extension policy to selected file systems.
8. In the **How often to run** area, schedule how often you want to enforce the policy. You can choose from the following actions:

**Enforce policy after every Probe or Scan**

Enforce the policy after every probe or scan that is run against the selected file systems. If you want to automate file system extension, choose to apply the policy after every new probe or scan.

**Run Now**

Enforce the policy immediately.

**Run Once at**

Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

**Note:** The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

**Note:** If the policy applies to multiple file systems and you select **Run Now** or **Run Once at**, messages specific to each file system are placed in the log file. Messages for multiple file systems can make it difficult to troubleshoot errors. To avoid this issue, you can select one policy at a time and then select **Run Now**.

9. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
  - a. 9:00 AM of the time zone where the Data server resides
  - b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

**Use the timezone that the server runs in**

Select this option to use the time zone of the location where the Data server resides.

**Use this timezone**

Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field is activated.

10. Click the **Alert** tab; the Alert page opens. Use this page to define the alerting criteria for a file system extension.
11. In the **Triggering Condition** area, specify the alerting criteria for the file system extension. In the **Condition** field, the default setting is **A filesystem extension action started automatically**. No other conditions are available.

When a file system is extended by the policy, an alert is triggered. When the alert is triggered, the actions you specified in the Triggered Actions section are taken.

12. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition. You can choose from the following actions:

**SNMP trap**

An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

**TEC / OMNibus Event**

A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, and so on.

**Login notification**

The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

**Login ID**

The ID of the user who receives alert notification upon logging in to the system.

**Windows Event Log, UNIX Syslog**

The ID of the user who receives alert notification upon logging in to the system.

**Event Type (Windows Event Log only)**

The type of event that is recorded to the OS log.

**Facility (UNIX Syslog only)**

The location where the UNIX Syslog is sent. You can select User or Local.

**Level (UNIX Syslog only)**

The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

**Run script**

A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window are passed to the script when the alert is triggered

**E-mail** An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit e-mail** to change the actual email message.


13. Save the file system extension:

- a. Click **File > Save**. The Save As window opens.
- b. In the **Specify Filesystem Extension Rule name** field, type a name for the job and click **OK**. The file system extension is displayed in the navigation tree pane.

## Editing a file system extension

Learn how to edit a file system extensions policy to change the settings for that policy.

To edit a file system extension policy, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager > Policy Management > Filesystem Extension**.
2. Click the file system extension you want to edit. Information about the file system extension you want to edit appears in the right content pane.
3. In the **Description** field, type a brief description of the job.
4. Select the file systems for the policy:
  - a. In the **Available** field, click the file systems you want to include.
  - b. Click . The file systems are moved to the **Current Selections** field.
5. Click the **Extension** tab; the Extension page opens. Use this page to edit how file systems will be extended. You can choose to extend file systems:
  - By a specific amount of space
  - By a percentage of file system capacity
  - By increasing free space to a specific size
  - By increasing free space to a percentage of file system capacity
  - By increasing capacity to a specific size

Select the option you want to use for this policy. You can also set a maximum capacity size if necessary. If you specify a maximum size, a file system extension job that would increase a file system beyond that size will fail.

**Note:** The actual amount of file system extension can vary, depending on your environment. For example, if 5GB physical partitions are set up for a volume group, its file systems will be extended in 5GB increments. If you specify a smaller amount, it will be rounded up to the physical partition size.

As another example, software striping and software RAID will claim part of the space created when file systems are extended. In this case, file system extension will appear to create slightly less space than specified in the policy.

If you are using LUN provisioning, you can specify the size of new LUNs by using the Create LUNs that are at least option on the Provisioning tab and one of the Amount to Extend options on the Extension tab. If the Amount to Extend value is greater than or equal to the preferred minimum LUN size, 1% more space will be added to the LUN to cover overhead requirements.

If you choose to extend file systems regardless of remaining free space, the policy options you specify will be applied to all selected file systems when the policy is enforced. To use policy to automatically manage file system capacity, you must do the following:

- Use the **When to Run** page to set the policy to run after every probe or scan
- Use the **Extension** page to specify an amount of free space to be used as a condition for extension. You can do this two ways:
  - Select **Make Freespace** and specify an amount. Then select Extend filesystems regardless of remaining free space. When free space falls below the amount you specify, the difference in space will be added. This option can be used to maintain file system free space at a specific level.

- Select either **Add** or **Make Freespace** and specify an amount. Then select **Extend filesystems** when free space is less than and specify an amount. The amount you specify will be used to trigger the **Amount to Extend** action you selected. This option can be used to add extra space to a file system if free space reaches a specified level. For example, if free space reaches 10% or less, it could be increased to 20%, or 5GB could be added to the total file system capacity.

**Note:** The following restrictions apply when you create or edit a file system extension policy:

- If you select both **Make Freespace** and **Extend filesystems when freespace is less than**, you must choose the same type of extension for each. If you specify a percentage amount for one, you must specify a percentage amount for the other. If you specify an amount expressed in megabytes for one, you must specify an amount in either megabytes or gigabytes for the other.
- If you select both **Make Freespace** and **Extend filesystems when freespace is less than**, the extension amount you specify for **Make Freespace** must be larger than the amount you specify for **Extend filesystems when freespace is less than**.
- If you select **Make Capacity**, you cannot select **Extend filesystems when freespace is less than**. Because the **Make Capacity** option extends capacity regardless of actual utilization, it is not intended to be used for the automated maintenance of file system capacity. You can use this option to extend file system capacity when necessary by selecting **Run Now** or **Run Once** at on the **When to Run** page.

You can also choose to have no file systems extended when the policy is applied. Instead, any policy actions that would have been performed are written to the log file. This feature, called LOG ONLY mode, can be used to preview the results of a policy over time before actually extending any file systems.

6. Click the **Provisioning** tab; the Provisioning page opens. Use the Provisioning tab to enable LUN provisioning and specify how it will be performed. Provisioning is an optional feature that is not required for file system extension. If you enable LUN provisioning, two things will happen when there is not enough space in a volume group to extend a file system:
  - A new LUN is created to add space to the volume group, according to the policy specified on the **Provisioning** tab.
  - The file system is extended within the volume group, according to the policy specified on the **Extension** tab.

For example, a volume group might first be extended by 100Mb, and then one or more file systems within it might be extended by 10MB each.

Data Manager relies on information stored by a CIM/OM to interact with TotalStorage Enterprise Storage Server subsystems. To implement provisioning, CIM/OM access information must be defined to Data Manager. See the CIM/OM Logins section in the Configuration and Getting Started Guide for more information. After CIM/OM access is configured, a Discovery job must be run to identify TotalStorage Enterprise Storage Server subsystems.

**Note:** LUNs can be provisioned for file system hosts running Solaris, but the hosts must be specially configured to avoid a reboot after provisioning. See the appendix on configuring LUN provisioning for Solaris in the IBM Tivoli

Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide for instructions.

If you choose to enable automatic LUN provisioning, you must choose a model and location for new LUNs. Data Manager models new LUNs on existing LUNs. Modeled characteristics include CIM/OM access parameters, as well as connectivity and RAID type. You must specify how far Data Manager will look for a model LUNs, as well as limit where it will create new LUNs on the TotalStorage Enterprise Storage Server.

**Note:** To perform LUN provisioning, at least one TotalStorage Enterprise Storage Server LUN must be currently assigned to the volume group of the file system you want to extend, or to the managed host associated with the file system you want to extend. Select one of the following options to limit the scope of LUN modeling:

- Model new LUNs on others in the volume group of the file system being extended. If you select this option, Data Manager will only search for a model within existing LUNs in the volume group of the file system being extended.
- Model new LUNs on others on the same host as the file system being extended. If you select this option, Data Manager will first search for a model within existing LUNs in the volume group of the file system being extended, before looking further on the same host.

New LUNs can be created in the same Storage Pool as the file system being extended, or in any Storage Pool in the TotalStorage Enterprise Storage Server of the file system being extended. Select one of the following options to limit where new LUNs will be created:

- Provision new LUNs within the same Storage Pool as the file system being extended. For TotalStorage Enterprise Storage Server, a Storage Pool is a Volume Space. If you are using RAID, a Volume Space is equal to one RAID rank, which consists of a Disk Group (8 physical disks).
- Provision new LUNs within the same Storage System as the file system being extended. If you select this option, new LUNs can be created in any Storage Pool in the TotalStorage Enterprise Storage Server. Data Manager will first attempt to provision a new LUN in the same Storage Pool as the file system being extended, before looking further in the same Storage System.

You can optionally specify a preferred minimum size for new LUNs by selecting the **Create LUNs that are at least** option. If you select this option, LUNs will not be created any smaller than the size you specify. If you do not specify a preferred minimum size, the Amount to Extend specified for the policy will be used. For TotalStorage Enterprise Storage Server subsystems, the default minimum LUN size increment is 100 MB. New LUNs will be rounded up to satisfy this requirement. If the Amount to Extend value is greater than or equal to the preferred minimum LUN size you specify, 1% more space will be added to the LUN to cover overhead requirements. See Table 24 on page 491.

**Note:** Data Manager will only provision new LUNs. Existing unassigned LUNs will not be reused. Similarly, if provisioning fails after LUNs have been created, the new LUNs will not be deleted or reused. This limitation protects against accidentally overwriting existing data that might be associated with currently unassigned LUNs.

Table 24 on page 491 shows the possible outcomes when a new LUN is provisioned.



Table 24. LUN Provisioning Results

| Description                                                                                                            | Amount to Extend | TotalStorage Enterprise Storage Server minimum for new LUNs | Create LUNs that are at least (preferred minimum LUN size) | Actual LUN size that will be created | Actual amount by which the file system will be extended |
|------------------------------------------------------------------------------------------------------------------------|------------------|-------------------------------------------------------------|------------------------------------------------------------|--------------------------------------|---------------------------------------------------------|
| Minimum LUN size exceeds amount to extend                                                                              | 20Mb             | 100 Mb                                                      | 200Mb                                                      | 200Mb                                | 20Mb                                                    |
| Amount to extend equals preferred minimum LUN size                                                                     | 200Mb            | 100 Mb                                                      | 200Mb                                                      | 202Mb (1% added for overhead)        | 200Mb                                                   |
| Minimum LUN size exceeds amount to extend                                                                              | 100Mb            | 100 Mb                                                      | 200Mb                                                      | 200Mb                                | 100Mb                                                   |
| Amount to extend exceeds preferred minimum LUN size, but both are below TotalStorage Enterprise Storage Server minimum | 75Mb             | 100 Mb                                                      | 50Mb                                                       | 100Mb                                | 75Mb                                                    |
| Minimum LUN size not specified                                                                                         | 200Mb            | 100 Mb                                                      | Not specified                                              | 202Mb (1% added for overhead)        | 200Mb                                                   |

**Note:** IBM TotalStorage Enterprise Storage Server and the TotalStorage Enterprise Storage Server CIM/OM use multiples of 1000 bytes to represent kilobytes, megabytes, and gigabytes. Data Manager uses multiples of 1,024 bytes. Data Manager automatically converts this difference when it discovers TotalStorage Enterprise Storage Server capacity and utilization. If you are using IBM Tivoli Storage TotalStorage Enterprise Storage Server Expert or TotalStorage Enterprise Storage Server Specialist, the storage information they show will differ from what is shown by Data Manager.

Due to AIX Volume Manager limitations on AIX Striped Volumes, provisioning of new LUNs cannot be done. If the policy definition for a striped file system indicates to Enable Automatic LUN Provisioning on the Provisioning tab, then any LUNs that are created will be added to the Volume Group but will not be usable for Extension. If provisioning is attempted on an AIX Striped Volume, then you will see the message: STA0301W: To prevent runaway condition, we have removed the file system file system\_name from the policy, please fix the problem and add the file system back to the policy.

To avoid this problem, make sure the **Enable Automatic LUN Provisioning** option is disabled for Filesystem Extension policy definitions that have striped file systems on AIX.

7. Click the **When to Run** tab; the When to Run page opens. Use this page to edit when to apply the file system extension policy to selected file systems.
8. In the **How often to run** area, schedule how often you want to enforce the policy. You can choose from the following actions:

**Enforce policy after every Probe or Scan**

Enforce the policy after every probe or scan that is run against the

selected file systems. If you want to automate file system extension, choose to apply the policy after every new probe or scan.

#### **Run Now**

Enforce the policy immediately.

#### **Run Once at**

Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

**Note:** The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

**Note:** If the policy applies to multiple file systems and you select **Run Now** or **Run Once at**, messages specific to each file system are placed in the log file. Messages for multiple file systems can make it difficult to troubleshoot errors. To avoid this issue, you can select one policy at a time and then select **Run Now**.

9. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
  - a. 9:00 AM of the time zone where the Data server resides
  - b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

#### **Use the timezone that the server runs in**

Select this option to use the time zone of the location where the Data server resides.

#### **Use this timezone**

Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field is activated.

10. Click the **Alert** tab; the Alert page opens. Use this page to edit the alerting criteria for a filesystem extension.
11. In the **Triggering Condition** area, specify the alerting criteria for the file system extension. In the **Condition** field, the default setting is **A filesystem extension action started automatically**. No other conditions are available. When a file system is extended by the policy, an alert is triggered. When the alert is triggered, the actions you specified in the Triggered Actions section are taken.
12. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition. You can choose from the following actions:

#### **SNMP trap**

An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

#### **TEC / OMNibus Event**

A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage

Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, and so on.

**Login notification**

The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

**Login ID**

The ID of the user who receives alert notification upon logging in to the system.

**Windows Event Log, UNIX Syslog**

The ID of the user who receives alert notification upon logging in to the system.

**Event Type (Windows Event Log only)**

The type of event that is recorded to the OS log.

**Facility (UNIX Syslog only)**

The location where the UNIX Syslog is sent. You can select User or Local.

**Level (UNIX Syslog only)**

The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

**Run script**

A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window are passed to the script when the alert is triggered

**E-mail** An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit e-mail** to change the actual email message.

13. Click **File** > **->Save** to save the filesystem extension.

## Scheduling script-based actions

Use scheduled actions to run user-defined scripts against selected computers and computer groups.

You can access the scheduled actions facility using either Data Manager or Data Manager for Chargeback. The actions that you schedule to run are based on user-defined scripts, allowing you to use any third-party tools for actions, such as recovery or provisioning. You can schedule scripts to run:

- Immediately
- Once at a specified time and date
- Repeatedly according to a schedule you specify

Select the option you want and then edit the time and date information accordingly. You can also select the time zone for the schedule you specify. You can either select the time zone that the Data server is in or any other time zone.

**Note:** Scheduled actions should not be confused with triggered actions. Triggered actions enable you to run scripts based upon conditions and alerts that are

detected within your storage environment. Scheduled actions enable you to run scripts according to the schedule that you define, without having to associate the script with an alert.

## Running scripts using scheduled actions


Use the scheduled action facility to run a script according to a schedule that you define. The scheduled action facility is available in Data Manager and Data Manager for Databases.

Use scheduled actions to run scripts against selected computers and computer groups. The scheduled actions facility is script-based, so you can use any third-party tools for actions, such as recovery, or provisioning. You can schedule scripts to run:

- Immediately
- Once at a specified time and date
- Repeatedly according to a schedule you specify

Scheduled actions are different from triggered actions. Use triggered actions to run scripts based upon conditions and alerts that are detected within your storage environment. Use scheduled actions to run scripts according to the schedule that you define, without having to associate the script with an alert.

To create a scheduled action, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager > Policy Management > Scheduled Actions** or **Data Manager for Databases > Policy Management > Scheduled Actions**.
2. Right-click **Scripts**, and click **Create Script**. The Create Script panel opens.
3. In the **Description** field, type a brief description of the job.
4. Select computers and computer groups against which to run the action:
  - a. In the **Available** field, click the computers and computer groups against which you want to run the action. You can select individual computers as well as computer groups that you have previously defined using the group facility. Cluster resource groups in a PowerHA or MSCS clustered environment are listed as computers.
  - b. Click . The computers and computer groups are moved to the **Current Selections** field.
5. Click the **Script Options** tab; the Script Options page opens. Use this page to select the script that you want to run against the selected computers and computer groups.
6. Select a script from **Script Name**. The scripts that appear in **Script Name** are stored in the `\<TPC_installation_directory>\scripts` directory on the Data server. The default installation directory where scripts are located is:
  - Windows: C:\Program Files\IBM\TPC\data\scripts
  - UNIX, Linux: /opt/IBM/TPC/data/scripts or /usr/IBM/TPC/data/scripts

**Note:** The scripts that appear in the Script Name list are stored on the computer where the Data server component is located. If you want to run a script that is stored on a Storage Resource agent or Data agent, you must type the name of that script in the Script Name field.

7. Click the **When to Run** tab; the When to Run page opens. Use the When to Run page to specify when the scheduled action is run.

8. In the **How often to run** area, schedule how often you want to run the script associated with the action. You can choose from the following options:

**Run Now**

Run the script immediately.

**Run Once at**

Run the script once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

**Note:** The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This process ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

**Run Repeatedly**

Run the script repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
- Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark is displayed next to the days that you select.

9. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
  - a. 9:00 AM of the time zone where the Data server resides
  - b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

**Use the timezone that the server runs in**

Select this option to use the time zone of the location where the Data server resides.

**Use this timezone**


Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field is activated.

10. Click the **Alert** tab; the Alert page opens. Use the Alert page to define the alerting criteria for a scheduled action.
11. In the **Triggering Condition** area, the only condition is **Script Failed**. An alert is automatically triggered if the scheduled action fails to run.
12. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see “Triggered actions for alerts” on page 1172.
13. Save the scheduled action:
  - a. Click **File > Save**. The Save As window opens.
  - b. In the **Specify Script name** field, type a name for the job and click **OK**. The scheduled action is displayed in the navigation tree pane.

## Editing a scheduled action

Edit a scheduled action to change: the computers and groups against which the action is run, the script that is run, when the script is run, and the alerting conditions if the action fails to run.

To edit a scheduled action, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager > Policy Management > Scheduled Actions > Scripts** or **Data Manager for Databases > Policy Management > Scheduled Actions > Scripts**.
2. Select the scheduled action that you want to edit. Information about the action appears in the right content pane.
3. In the **Description** field, type a brief description of the job.
4. Select computers and computer groups against which to run the action:
  - a. In the **Available** field, click the computers and computer groups against which you want to run the action. Select individual computers, as well as computer groups that you have previously defined using the group facility. Note that cluster resource groups in a PowerHA or MSCS clustered environment are listed as computers.
  - b. Click . The computers and computer groups are moved to the **Current Selections** field.
5. Click the **Script Options** tab; the Script Options page opens. Use this page to change the script you want to run against the selected computers and computer groups.
6. Select a script from **Script Name**. The scripts that appear in **Script Name** are stored in the `\<TPC_installation_directory>\scripts` directory on the Data server. The default installation directory where scripts are located is:
  - Windows: C:\Program Files\IBM\TPC\data\scripts
  - UNIX, Linux: /opt/IBM/TPC/data/scripts or /usr/IBM/TPC/data/scripts

**Note:** The scripts that appear in the Script Name list are stored on the computer where the Data server component is located. If you want to run a script that is stored on a Storage Resource agent or Data agent, you must type the name of that script in the Script Name field.

7. Click the **When to Run** tab; the When to Run page opens. Use this page to change when the scheduled action is set to run.
8. In the **How often to run** area, schedule how often you want to run the script associated with the action. You can choose from the following options:

### Run Now

Run the script immediately.

### Run Once at

Run the script once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

**Note:** The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This process ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

### **Run Repeatedly**

Run the script repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
- Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark is displayed next to the days that you select.

9. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
  - a. 9:00 AM of the time zone where the Data server resides
  - b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

### **Use the timezone that the server runs in**

Select this option to use the time zone of the location where the Data server resides.

### **Use this timezone**

Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field is activated.

10. Click the **Alert** tab; the Alert page opens. Use the Alert tab to change the alerting criteria for a scheduled action.
11. In the **Triggering Condition** area, the only condition is **Script Failed**. This indicates that an alert will automatically be triggered if the scheduled action fails to run.
12. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition. You can choose from the following actions:

### **SNMP trap**

An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

### **TEC / OMNIBus Event**

A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It also performs responses such as sending further notification, creating or updating trouble tickets, running programs, and so on.

### **Login notification**

The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

### **Login ID**

The ID of the user who receives alert notification upon logging in to the system.

**Windows Event Log, UNIX Syslog**

The ID of the user who receives alert notification upon logging in to the system.

**Event Type (Windows Event Log only)**

The type of event that is recorded to the OS log.

**Facility (UNIX Syslog only)**

The location where the UNIX Syslog is sent. You can select User or Local.

**Level (UNIX Syslog only)**

The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

**Run script**

A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window are passed to the script when the alert is triggered

**E-mail** An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit e-mail** to change the actual email message.

13. Click **File > Save** to save your changes.

## Archive and backup functions

Data Manager provides a method within the alerting function of a constraint by which you can automatically invoke a IBM Tivoli Storage Manager archive or backup to run against the files that violate that constraint.

An archive is a "snapshot in time". Each time an archive is done, a new "snapshot" is taken. When an archive is performed, everything that has been specified is captured. So, for example, if an archive was performed on a C: drive, everything on C: will be archived. If, the next day, another archive of C: was performed, everything on the C: drive would be archived again. Files and file systems that have been archived as part of the Archive/Backup job can be deleted after archiving, but you must select **Delete After Successful Archival** specifically.

Lets say that, as a result of running a report, you find a large number of files that haven't been accessed in over two years. Looking at the cumulative sizes of these files you determine they are taking up space that would be better utilized elsewhere. You can archive and delete these files to reclaim that space. If, at some time in the future those files are needed, they can be retrieved from the archive. The result of the archive job is more free space without data loss.

A backup can be incremental (**Selective Backup**) or full. Usually a full backup is done first and subsequent backups are incremental. For example, say that a backup is performed on C: drive. All the files on the C: drive are then backed up, which is the full backup operation. The next day we need to back up C: drive again. Usually, this will be an incremental backup since a full backup was done the day before. In the incremental backup, only the files that have changed on the drive since the last backup will be backed up.

The archived or backed up files are sent to theTivoli Storage Manager server and are stored on that machine. Files are put into storage groups that are controlled through a Tivoli Storage Manager policy that the Tivoli Storage Manager server administrator usually creates.



## Constraints

Data Manager provides a method within the alerting function of a constraint by which you can automatically invoke a Tivoli Storage Manager archive or backup to run against files that violate constraints you have defined. You can do the following:

- Define a constraint that will invoke Tivoli Storage Manager to archive and (if that option is selected) delete those files
- Define a constraint to run an incremental or full backup

See the scenario *Archive and delete files* to see an example of how constraints are developed.

## Before you begin

Before you can use the Tivoli Storage Manager archive/backup functionality with Data Manager constraints, you must ensure the following:

- Your user ID has Administrator authority within IBM Tivoli Storage Productivity Center.
- Tivoli Storage Manager client software is installed on each of the hosts where a Storage Resource agent or Data agent is used to perform archive and backup functions.
- (for SAN attached storage only) Tivoli Storage Manager client software must be installed on the Storage Resource agent or Data agent host computer that is configured to act as the scanning agent for NAS storage.
- A Tivoli Storage Manager client must be configured to run without a password. This is typically done by either including the password in the client options (dsm.opt) file, or using the options `PASSWORDACCESS= GENERATE` in the client options file.
- The Tivoli Storage Manager client must be version 5.1 or higher.

## Create an archive/backup job

Create a IBM Tivoli Storage Manager archive/backup job for files in a report

You follow these general steps when you create an archive or backup.

1. Make sure you have followed the prerequisite steps
2. Decide whether to create a new constraint or use one of the default constraints
3. Generate the report using the constraint you decided to use
4. Create an archive or backup job from files in the report

To create an Archive/Backup job on files from a report, complete the following steps:

1. Generate and view a Largest Files, Most Obsolete Files, Orphan Files, Duplicate Files, Most At Risk, Constraint Violations report, or a report that you have created.
2. Click the files you want to include in the archive or backup job. Use ctrl + click to select multiple, non-consecutive files; use shift + click to select multiple, consecutive files.
3. Right click the reports and select **Create a new archive/backup job** from the pop-up menu. The **Create File Archive/Backup Action** window is displayed.
4. In the **Description** field, type a brief description of the job.

5. In the **Archive/Backup Options** tab, select the type of the job and its TSM command options, if any:

#### **Archive**

Use this option to archive the files that violate a constraint. Archives are copies of data from online storage copies to offline storage.

#### **Delete After Successful Archival**

Automatically delete the files that violated the constraint after they have been successfully archived. See the *IBM Tivoli Storage Manager Client User's Guide* for more information about Tivoli Storage Productivity Center archives.

**Note:** Tivoli Storage Productivity Center administrators (Productivity Center Administrator role) have the ability to delete files even if they do not have the required authority on the systems where the files are stored. This authority is provided by the Tivoli Storage Productivity Center agents which run as root on UNIX systems or with administrator privileges on Windows computers.

#### **Selective Backup**

Select this option to perform a selective backup of the files that violate a constraint. A selective backup enables you to back up specified files. These files are not excluded in the include-exclude list and meet the requirement for serialization in the backup copy group of the management class assigned to each file. See the *IBM Tivoli Storage Manager Client User's Guide* for more information about selective backups.

#### **Incremental Backup**

Select this option to perform an incremental backup of the files that violate a constraint. An incremental backup enables you to back up new or changed files or directories from a client domain or from specified directories or files. These directories or files are not excluded in the include-exclude list and meet the requirements for frequency, mode, and serialization as defined by a backup copy group of the management class assigned to each file. See the *IBM Tivoli Storage Manager Client User's Guide* for more information about incremental backups.

#### **IBM Tivoli Storage Manager Command Options**

Enter additional Tivoli Storage Manager command line options in this field that will be passed on to the Tivoli Storage Manager `dsmc` command. You can enter up to 512 bytes of text. See the *IBM Tivoli Storage Manager Client User's Guide* for more information about valid Tivoli Storage Manager command-line options.

6. Click the **Files** tab; the **Files** page opens. Use this page to view or remove the files included in an archive/backup job.
7. Click the **When to Run** tab. Use the **How often to run** area to specify when the archive or backup job runs.
8. In the **How often to run** area, schedule how often you want the quota to be checked. You can choose from the following actions:

#### **Run Now**

Check the quota immediately.

### Run Once at

Check the quota once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

**Note:** The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

### Run Repeatedly

Check the quota repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
  - Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
  - Indicate which days on which run the job. A check mark is displayed next to the days you select.
9. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
- a. 9:00 AM of the time zone where the Data server resides
  - b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

### Use the timezone that the server runs in

Select this option to use the time zone of the location where the Data server resides.

### Use this timezone

Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field is activated.

10. Click the **Alert** tab; the Alert page opens. Use the Alert page to define the alerting criteria for an archive or backup job.
- a. In the **Triggering Condition** area, the only condition is **File Archive/Backup Failed**. This condition indicates that an alert is triggered automatically if the archive or backup fails to run.
  - b. In the **Triggered Actions** area, select the actions that occur as a result of the triggering condition You can choose from the following actions:

#### SNMP trap

An SNMP trap is sent to any NMS, console, or terminal when the condition occurs.

#### TEC / OMNibus Event

A IBM Tivoli Enterprise Console event is sent when the condition occurs. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze Tivoli Storage Productivity Center events according to the event definitions specified in the tivoliSRM.baroc file (provided with the product). It

also performs responses such as sending further notification, creating or updating trouble tickets, running programs, and so on.

**Login notification**

The alert is sent to a Tivoli Storage Productivity Center user. The user receives the alert upon logging in to Tivoli Storage Productivity Center. In the **Login ID** field, type the user ID.

**Login ID**

The ID of the user who receives alert notification upon logging on to the system.

**Windows Event Log, UNIX Syslog**

The ID of the user who receives alert notification upon logging on to the system.

**Event Type (Windows Event Log only)**

The type of event that is recorded to the OS log.

**Facility (UNIX Syslog only)**

The location where the UNIX Syslog is sent. You can select User or Local.

**Level (UNIX Syslog only)**

The level of UNIX Syslog event. You can select Informational, Notice, Warning, Error, Critical, or Alert.

**Run script**

A script is run when the condition occurs. Click **Define** to specify the script that is run when the triggering condition occurs. The parameters listed in the **Specify Script** window are passed to the script when the alert is triggered

**E-mail** An email is sent. Click **Add** or **Delete** to edit the list of email addresses to which an email is sent. Click **Edit e-mail** to change the actual email message.

11. Save the archive or backup:
  - a. Click **File > Save**. The **Save As** window opens.
  - b. In the **Specify File Archive/Backup Action name** field, type a name for the job and click **OK**. The new job is listed in **Data Manager > Policy Management > Archive/Backup** as **<creator>.<job name>**.

**View and edit an Archive/Backup job**

Learn how to edit IBM Tivoli Storage Manager archive and backup jobs that were defined to run against the files selected from Data Manager reports or from reports you have generated yourself.

Each archive and backup that you have created appears as a node under **Archive/Backup**.

To view and edit a Tivoli Storage Manager archive or backup job, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager > Policy Management > Archive/Backup**. Select the archive or backup job that you want to edit. Information about the job appears in the right content pane. All the user-defined information can be edited.
2. Change the description for the archive/backup job.

3. In the **Options** tab, you can select the Archive/Backup options. You can change whether the job is an archive or a backup, and add or change any additional Tivoli Storage Manager commands that should be run with the job.
4. In the **Files** tab the files that were included in the job are listed.
5. In the **When to Run** tab, select how often to run and how to handle time zones. You can change the **How often to run** and **How to handle time zones** settings.
6. In the **Alert** tab, **File Archive/Backup Failed** is the only **Triggering-Condition**. In **Triggered-Actions**, change the action you wish to occur if the job fails. You do not need to select any of the actions. Notification will be written to the error log by default.
7. Select **File > Save**. Name the Archive/Backup job. After you save a the job, an entry for it will appear under the **Data Manager > Policy Management > Archive/Backup > <Creator>.<Name>**.

### **Add additional files to an archive/backup job**

Add files to an existing IBM Tivoli Storage Manager archive or backup job from a report.

To add files to an existing Archive/Backup job, complete the following steps:

1. Add files to an existing Tivoli Storage Manager archive or backup job from a report
2. Click the files that you want to add to an archive or backup job. Use ctrl + click to select multiple, non-consecutive files; use shift + click to select multiple, consecutive files.
3. Right click the report and select **Add to an existing archive/backup job** from the pop-up menu. The Archive/Backup Job Selection window appears.
4. Select the archive or backup job to which you want to add the files from the list box.
5. Click **OK** to add the files. The Archive/Backup - Files window appears and shows the file(s) you added to the selected job.

---

## **Working with storage resource groups**

Use storage resource groups to organize logically related storage entities into named groups. For example, a storage resource group that represents a business critical email application might include all the hosts that participate in the email application cluster, the storage subsystems that provide storage to the application, and the switches and fabrics through which the application's data is configured to travel.

You can use storage resource groups as input to the SAN Planner for planning and provisioning considerations. For example, if a storage resource group represents a business critical email application, you can use that group as input into the SAN planner to help consistently determine when and where to add capacity to the application while minimizing the manual entry of storage entities.

### **Storage resource group management**

Use the storage resource group management list panel to view, create, and delete storage resource groups.

This panel contains the following information about each storage resource group:

**Storage resource group name**

Specifies a user-defined name of the storage resource group.

**State** Specifies the state of the storage resource group:

**Normal**

The storage resource group is operating normally.

**Warning**

At least one part of the storage resource group is not operating or has serious problems.

**Critical**

The storage resource group is either not operating or has serious problems.

**Missing**

The storage resource group was previously discovered by IBM Tivoli Storage Productivity Center, but was not detected in the most recent discovery.

**Unknown**

The storage resource group was discovered but is not recognized by IBM Tivoli Storage Productivity Center.

**Note:** The state of a storage resource group might appear as **Unknown** after it is created. This occurs because the status of a new storage resource group might not yet be updated when the Storage Resource Group Management panel is re-displayed. To resolve this issue, exit the Storage Resource Group Management panel, wait for the processing of the group to complete, and open the Storage Resource Group Management panel again to view the updated **State**.

**Description (optional)**

Displays the user defined description for the storage resource group.

**Create** Creates a storage resource group that you selected.

**Delete**

Deletes a storage resource group that you selected.

## Create or Edit a Storage Resource Group

Use the Create or Edit Storage Resource group panel to create or modify storage resource groups.

To create a storage resource group:

- In the navigation tree, expand **IBM Tivoli Storage Productivity Center**, click **Storage Resource Group Management**, then click **Create...** in the content pane.
- In the navigation tree, expand **IBM Tivoli Storage Productivity Center > Topology**, click a node that represents a storage entity, right-click an entity that appears in the content pane, and select **Add to new Storage Resource Group...** from the drop down menu.

To edit a storage resource group:

- In the navigation tree, expand **IBM Tivoli Storage Productivity Center**, click **Storage Resource Group Management**, then click the magnifying glass icon next to the storage resource group you want to edit.

- In the navigation tree, expand **IBM Tivoli Storage Productivity Center > Topology**, click **Storage Resource Groups**, right-click the storage resource group you want to edit in the content pane, and select **Launch Detail Panel** from the drop down menu.

### Field Descriptions

#### Creator

Displays the user name of the creator.

**Name** Displays the name of the storage resource group or unnamed, if it is not yet named.

#### Description

Optional: Displays the user defined description for the storage resource group.

#### Selected Elements

Lists the elements selected to be members of this storage resource group.

**Add** Adds one or more selected elements to the list.

The Storage resource group element selection panel is displayed.

#### Remove

Removes one or more selected elements from the list.

#### Default Provisioning Profile

Lists the available provisioning profiles which can be associated with storage resource groups. The list also includes None. If this storage resource group is used as input to the SAN Planner, the settings defined in this profile will be used to pre-populate the planner inputs.

#### Create a New Profile

Launches the Provisioning Profile creation wizard. When you complete the wizard, the **Provisioning Profile** list is updated.

#### User defined property 1 (UDP1)

Specifies any user-defined properties that will be used by the topology viewer to provide custom groupings.

#### User defined property 2 (UDP2)

Specifies any user-defined properties that will be used by the topology viewer to provide custom groupings.

#### User defined property 3 (UDP3)

Specifies any user-defined properties that will be used by the topology viewer to provide custom groupings.

### Storage Resource Group Element Selection

Use the Storage Resource Group Element Selection panel to select the storage entities that you want to include in a storage resource group. The storage entities that you can select are displayed in the topology viewer.

### Available elements section

Displays the storage entities that you can include in a storage resource group. Use the topology viewer to select storage entities. See Overview of Topology Viewer for more information about the topology viewer.

### Topology viewer

Use the topology viewer to select the storage entities that you want to include in a storage resource group. To do this, perform the following steps.

1. Open a detailed view of a storage entity type on the topology viewer's **Overview** page.
2. Expand the storage entity type to view the storage entities that are part of that type.
3. Select the storage entities you want to include in the storage resource group.
4. Click >>.

### Selected elements section

>> Click >> to move the selected subsystems into the **Selected Elements** pane.

<< Click << to remove the selected subsystems from the **Selected Elements** pane.

### Selected Elements

Displays the storage entities that are included in a storage resource group. You can include the following entities in a storage resource group:

- computer (vm, physical or hypervisor)
- computer disk (including vm disks)
- fabric
- switch
- other entities
- storage subsystem
- storage pool
- subsystem disk
- subsystem volume
- tape drive
- tape library
- tape media changer
- other storage resource groups

**OK** Click **OK** when you are satisfied with all your selections. The Create Storage Resource Group panel displays with your selections.

### Cancel

Click **Cancel** if you do not want to add storage entities to a storage resource group. The Create Storage Resource Group panel displays. No data is saved.

### Working with storage resource group sessions

This topic describes viewing and deleting replication sessions that are created in the SAN Planner wizard for a storage resource group.

You can delete the replication session from IBM Tivoli Storage Productivity Center and optionally also from IBM Tivoli Storage Productivity Center for Replication.

1. In the left pane, expand **Replication Manager > Storage Resource Group Sessions**. A table of the currently defined storage resource group sessions is displayed in the right pane.



2. To select a storage resource group session to delete, click a row in the table. You can use Ctrl+click and Shift+click to select multiple sessions simultaneously.
3. Click **Delete Storage Resource Group Session**. Tivoli Storage Productivity Center does not automatically delete the replication session from Tivoli Storage Productivity Center for Replication. A dialog box opens enabling you to specify how you want to delete the session. You can select from the following options:
  - Click **Yes** to remove the selected replication session from the Tivoli Storage Productivity Center database and the Tivoli Storage Productivity Center for Replication database.
  - Click **No** to remove the selected replication session from the Tivoli Storage Productivity Center database only.
  - Click **Cancel** to close the dialog box without deleting the selected replication session from the Tivoli Storage Productivity Center database or Tivoli Storage Productivity Center for Replication database.

---

## Working with scripts

Use the triggered action facility to run scripts based upon conditions/alerts that are detected within your storage environment. Use the scheduled action facility to run scripts according to the schedule that you define, without having to associate the script with an alert.

Both the triggered action and scheduled action facilities are script-based, allowing you to use any third-party tools for actions, such as recovery and provisioning.

## Running scripts with alerts

Monitoring, alerting, and policy management jobs enable you to use the triggered action facility to run scripts based on storage events that you define.

The triggered action facility allows you to define thresholds and constraints that result in notification or action. A script can be run when a job fails, an alert condition is detected in a storage resource, a constraint is violated (Data Manager only), or a quota is exceeded (Data Manager only). The triggered action facility is script-based, allowing you to use any third-party tools for actions such as recovery or provisioning.

Perform the following steps to run a script in response to a condition in a job:

1. Create or edit a monitoring, alerting, or policy management job. For example, the following jobs enable you to define a condition that triggers a script:
  - Data Manager:
    - **Monitoring** > <all\_monitoring\_jobs>
    - **Alerting** > <all\_alert\_jobs>
    - **Policy Management** > **Constraints**
    - **Policy Management** > **Quotas**
    - **Policy Management** > **Scheduled Actions**
  - Data Manager for Databases:
    - **Monitoring** > <all\_monitoring\_jobs>
    - **Alerting** > <all\_alert\_jobs>
    - **Policy Management** > **Quotas**
    - **Policy Management** > **Scheduled Actions**

- Disk Manager:
    - **Monitoring > Subsystem Performance Monitors**
    - **Alerting > Storage Subsystem Alerts**
  - Fabric Manager:
    - **Monitoring > Switch Performance Monitors**
    - **Alerting > <all\_alert\_jobs>**
2. Check **Run Script** on the **Alert** page for the job.
  3. Click **Define**. The **Specify Script** dialog box displays.
  4. Enter the name of the script in the **Script Name** field. This script name is used to tell the server what script to send to the agent. If the script sent by the server is run by the agent, the name of the script will not be the name you specify in the **Script Name** field. Instead, the agent names the script as temporary file, runs it, then deletes it after the script has run. Because of this, you should not have code in your script that depends on this script name.
  5. Specify where you want the script to run in the **Where to Run** field. You can select a specific agent, or the triggering computer. For example, if a file system alert is triggered on a on a specific computer, you would typically want the script to run on the computer that triggered the alert. (For some types of alerts such as *Job Failed*, you will not have the option to run the script on the triggering computer. Here you can run the script on any agent you choose.)
  6. The **Script Parameters** section displays what parameters will be passed to the script when it is run. Each triggering condition passes different parameters. The parameters will be passed in the order specified on the dialog box.
  7. Click **OK**.

## Running scripts using scheduled actions

Use the scheduled action facility to run a script according to a schedule that you define. The scheduled action facility is available in Data Manager and Data Manager for Databases.


Use scheduled actions to run scripts against selected computers and computer groups. The scheduled actions facility is script-based, so you can use any third-party tools for actions, such as recovery, or provisioning. You can schedule scripts to run:

- Immediately
- Once at a specified time and date
- Repeatedly according to a schedule you specify

Scheduled actions are different from triggered actions. Use triggered actions to run scripts based upon conditions and alerts that are detected within your storage environment. Use scheduled actions to run scripts according to the schedule that you define, without having to associate the script with an alert.

To create a scheduled action, complete the following steps:

1. In the Navigation tree pane, expand **Data Manager > Policy Management > Scheduled Actions** or **Data Manager for Databases > Policy Management > Scheduled Actions**.
2. Right-click **Scripts**, and click **Create Script**. The Create Script panel opens.
3. In the **Description** field, type a brief description of the job.
4. Select computers and computer groups against which to run the action:

- a. In the **Available** field, click the computers and computer groups against which you want to run the action. You can select individual computers as well as computer groups that you have previously defined using the group facility. Cluster resource groups in a PowerHA or MSCS clustered environment are listed as computers.
  - b. Click . The computers and computer groups are moved to the **Current Selections** field.
5. Click the **Script Options** tab; the Script Options page opens. Use this page to select the script that you want to run against the selected computers and computer groups.
  6. Select a script from **Script Name**. The scripts that appear in **Script Name** are stored in the `\<TPC_installation_directory>\scripts` directory on the Data server. The default installation directory where scripts are located is:
    - Windows: C:\Program Files\IBM\TPC\data\scripts
    - UNIX, Linux: /opt/IBM/TPC/data/scripts or /usr/IBM/TPC/data/scripts

**Note:** The scripts that appear in the Script Name list are stored on the computer where the Data server component is located. If you want to run a script that is stored on a Storage Resource agent or Data agent, you must type the name of that script in the Script Name field.

7. Click the **When to Run** tab; the When to Run page opens. Use the When to Run page to specify when the scheduled action is run.
8. In the **How often to run** area, schedule how often you want to run the script associated with the action. You can choose from the following options:

**Run Now**

Run the script immediately.

**Run Once at**

Run the script once at a specified time and date. You can use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM.

**Note:** The Tivoli Storage Productivity Center framework that starts the Tivoli Storage Productivity Center jobs finds all jobs in an interval between one minute to the scheduled time and the scheduled start time. This process ensures that if many jobs are planned to run at a certain time, they will not all start after the scheduled start time.

**Run Repeatedly**

Run the script repeatedly according to a schedule you specify. You can:

- Use the drop-down lists to specify the month, day year, and the hour, minute, and AM/PM when the job should begin running.
- Indicate how often the job should run. You can use the drop-down lists to specify the number of minutes, hours, days, weeks, or months.
- Indicate which days on which run the job. A check mark is displayed next to the days that you select.

9. In the **How to Handle Time Zones** area, indicate the time zone to be used when scheduling and running the policy. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run it at:
  - a. 9:00 AM of the time zone where the Data server resides

- b. 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the job will run at 7:00 AM on an agent located in the PST time zone.

**Use the timezone that the server runs in**

Select this option to use the time zone of the location where the Data server resides.

**Use this timezone**

Select this option to indicate a specific time zone for when to run the job. When you select this option, the drop-down list to the right of the field is activated.

10. Click the **Alert** tab; the Alert page opens. Use the Alert page to define the alerting criteria for a scheduled action.
11. In the **Triggering Condition** area, the only condition is **Script Failed**. An alert is automatically triggered if the scheduled action fails to run.
12. In the **Triggered-Actions** pane, select the actions that occur as a result of the triggering condition. For an explanation of triggered actions, see “Triggered actions for alerts” on page 1172.
13. Save the scheduled action:
  - a. Click **File > Save**. The Save As window opens.
  - b. In the **Specify Script name** field, type a name for the job and click **OK**. The scheduled action is displayed in the navigation tree pane.

## How scripts are run

Scripts can be run from the Data server, Data agents, or Tivoli Storage Productivity Centers.

During installation, IBM Tivoli Storage Productivity Center creates a directory on every server and agent computer called: `<install_directory>\scripts`, where `<install_directory>` represents the directory where the agent or server is installed. The default installation directory is:

- (Windows Data agents) `C:\Program Files\IBM\TPC\ca\subagents\TPC\Data\scripts`
- (UNIX, Linux Data agents) `/opt/IBM/TPC/ca/subagents/TPC/Data/scripts` or `/usr/tivoli/ep/subagents/TPC/Data/scripts`
- (Windows Storage Resource agents) `C:\Program Files\IBM\TPC\agent\scripts`
- (UNIX, Linux Storage Resource agents) `/opt/IBM/TPC/agent/scripts` or `/usr/tivoli/ep/subagents/TPC/Data/scripts`
- (Windows Data server) `C:\Program Files\IBM\TPC\Data\scripts`
- (UNIX, Linux Data server) `/opt/IBM/TPC/Data/scripts` or `/usr/IBM/TPC/Data/scripts`

When a script is run:

1. The Agent looks in the agent's local `\scripts` directory; if the Agent can resolve that script name, it is executed from that location.
2. Otherwise, (if the agent cannot resolve the script name) and the Agent is configured to accept scripts from the server, the server looks in the local `\scripts` directory. If the server can resolve that script name, it loads that script and sends it to the agent where you designated the script to be run.

When running a script against a NAS filer, the script is run from the agent assigned to the file system where the triggering condition occurred.

3. The Agent receives the script, names it as a temporary file, and runs it.
4. After the Agent has finished running the script, the temporary file is deleted.  
Because the script is handled in this way, any script which is written to redirect any output must specify a fully-qualified unique directory path for that output.

**There are two main reasons to run a script:**

- In response to an alert
- For batch reports

**There are two reasons why a script might not run when sent by the server:**

- The script exists in an agent's \scripts directory. In this case, the agent ignores the script sent by the server and run its local version.
- If you did not check the **Agent may run scripts sent by server** option in the installation parameters screen, the agent ignores the script sent by the server.

If you selected the **Agent may run scripts sent by server** option in the installation parameters screen, the agent runs scripts sent by the server.

The advantage of selecting this option is you need only store one copy of the script in the \scripts directory for the server and do not have to keep a copy of the script on every agent computer. When a script must be run on a particular agent, the server accesses the script from the local \scripts directory and sends it to the agent to be run.

**Note:**

- You can change this option by editing the honorSentScripts parameter in the agent.config file. The agent.config file exists in every installation directory for the agent. See *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* for information about the agent.config file.
- After a Storage Resource agent is deployed, you can change the script configuration in the user interface. In the navigation tree, expand **Administrative Services > Data Sources > Data/Storage Resource Agents**. Select the target Storage Resource agent from the list, then open the **Select Action** menu and select the **Enable Scripts** or **Disable Scripts** option as appropriate.

When the script is deployed to the agent, it is named as a temporary file. After the agent has finished running the script, the temporary file is deleted.

If the **Agent may run scripts sent by server** option is cleared, you must make sure that the script is stored in every \scripts directory for the agent. If the **Agent may run scripts sent by server** option is checked, the server deploys the script to the agent computer that has the low free space condition. In this case, you must ensure that the script exists in the \scripts directory for the server, and not on any of the \scripts directory for the agents.

**Note:**

- A good practice is to store the scripts you want to run globally (across all your agents) on the server. Store the scripts you want to run locally on the agent computers where you want them to run.
- When running a script for Data agents (4.1.1 and lower), the job is immediately indicated as successfully run. Tivoli Storage Productivity Center does not wait for the script results. If there are any script errors, the errors are not reflected in the job status.

## How the server and agent resolve script names

An IBM Tivoli Storage Productivity Center server processes requests for agents running under UNIX, Linux, and Windows.

When associating and running scripts with triggered and scheduled actions, the Tivoli Storage Productivity Center server resolves the names of those scripts based on the following criteria:

- The OS under which an agent is running
- How the script name is specified on the Specify Script dialog box

### Agent and Operating System

The operating system under which an agent runs determines the scripts that can be run by that agent:

- **UNIX, Linux.** An agent running under a UNIX or Linux operating system will not run or receive (from the server) scripts that have an extension. Agents running under UNIX or Linux only accept and run script files that do not contain an extension within their name. The first line in a script file determines what interpreter is needed to run the script.
- **Windows.** An agent running under a Windows operating system will only run or receive (from the server) scripts that have an extension. The extension of a script file determines what interpreter is needed to run the script.

### Specifying a Script Name

Use the Specify Script dialog box when associating a script with a triggered or scheduled action. Enter the name of a script on the **Script Name** field to associate that script with the action you want to perform. The name of the script you enter might or might not contain an extension (including an extension is optional). Depending on the agent that receives the script, the following occurs after a triggered or scheduled action that contains a script is run:

- **Agent running under UNIX or Linux.** If the specified script name contains an extension, the agent ignores the extension and searches for a file of the same name (without an extension). For example, if you enter backup.vbs in the **Script Name** field, a UNIX or Linux agent will search for a file with the name backup. If the script exists, UNIX or Linux will read the first line of the script and use the appropriate interpreter to run the script. If the file does not exist, the triggered or scheduled action will not run the script.
- **Agent running under Windows.** If the specified script name contains an extension, the agent will search for a file of the same name/extension. For example, if you enter backup.vbs in the **Script Name** field, an agent will search for a file with the name backup.vbs. If the script does exist, the agent will determine what interpreter is needed to run the script based on its extension. If the file does not exist, the triggered or scheduled action will not run the script.  
If the specified script name does not contain an extension, an agent will look for the file name in the following order (based on its extension): .com, .exe, .bat, .cmd, .vbs, .vbe, .js, .jse

The process for script name resolving enables you to run scripts with the same name across multiple operating systems. To run a script on both a UNIX or Linux and Windows operating systems, you must have two versions of that file on either the server or the agent. For example, if you want to run a provisioning script

against a UNIX or Linux and Windows agent, you must have two versions of that file (provision.bat and provision) stored in the \scripts directory on either the server or agent.

---

## Using the topology viewer

This topic describes how to use the topology viewer to view detailed information about your storage environment and monitor entities and groups of entities.

You can view health status, performance, and other details about switches, computers, storage subsystem servers, fabrics, and other entities. To launch the topology viewer, expand **IBM Tivoli Storage Productivity Center Topology** in the Navigation Tree pane.

To display detailed information about an entity (for example, the host name of a computer, the used and unused capacity of a storage subsystem), right-click the entity and select **Launch Detail Panel**.

To select an entity in the graphical view, click the entity image. In the table view, click the name of the entity. To make multiple selections, press CTRL and click.

### Note:

- The availability of certain action items depends on the entity class and view level.
- When you right-click an object in the graphical view, a context menu is displayed. The action items in this menu are identical to the action items displayed when you right-click the corresponding object in the table view and click the arrow in the **Action** field.
- If the discovery process is not completed and there are entities in the environment that have not been rendered into the current view, the view will build dynamically as entities are discovered.
- Some columns and fields displayed within the Topology Viewer might not be populated with data for the storage resources that you are monitoring. Columns and fields appear empty if you have not configured Tivoli Storage Productivity Center to discover or collect the information intended for those columns and fields. Use discovery, probe, ping, or scan schedules to ensure that you collect information for all the columns and fields that appear in the Topology Viewer.

## The topology viewer

This topic describes how the topology viewer provides a graphical representation of the physical and logical resources that have been discovered in your storage environment by IBM Tivoli Storage Productivity Center.

The information displayed by the topology viewer is collected by discovery, probing, performance data collection, and alerting. The topology viewer consists of two views, a graphical view and a table view, organized vertically with the graphical view on top of the table view. The table view shows the same information as the graphical view, but in a tabular format.

The topology viewer supports four levels of detail, from a high-level overview of the entire storage environment to detailed information about individual entities. This allows you to view as much or as little detail as you need.

## Overview

A global, highly aggregated view of the entire storage environment. This is the default view when you expand **IBM Tivoli Storage Productivity Center > Topology**.

## Groups level (L0)

This level focuses on groups of a particular type. The default groups are:

- **Computers:** All discovered computer groups in your storage environment
- **Fabrics:** All fabric groups in your storage environment. A *fabric* is a network of entities that are attached through one or more switches. The topology viewer also displays Virtual SANs (VSANs) and Logical SANs (LSANs).
- **Switches:** All switches in your environment.
- **Storage:** All storage subsystems and tape libraries in your environment.
- **Other:** All entities that are discovered in your environment but that do not fit in either the computer, fabric or storage classes.

The groups level is the default level upon launching the topology viewer by clicking either **Computers**, **Fabrics**, **Switches**, **Storage** or **Other** under the **Topology** node. The groups level shows one or more groups of entities that correspond to the topology class that you selected. Information in the graphical view at this level can include any individually pinned entities.

## Group level (L1)

This level focuses on one selected group and its related entities (for example, a group of computers). At L1, individual entities can be shown by expanding the group box. You can also display lines representing connections between entities or between entities and groups.

## Detail level (L2)

This level focuses on individual entities (for example, a single computer, switch, subsystem, or tape library) and the paths to associated logical and physical entities. This level shows details all the way to the port level.



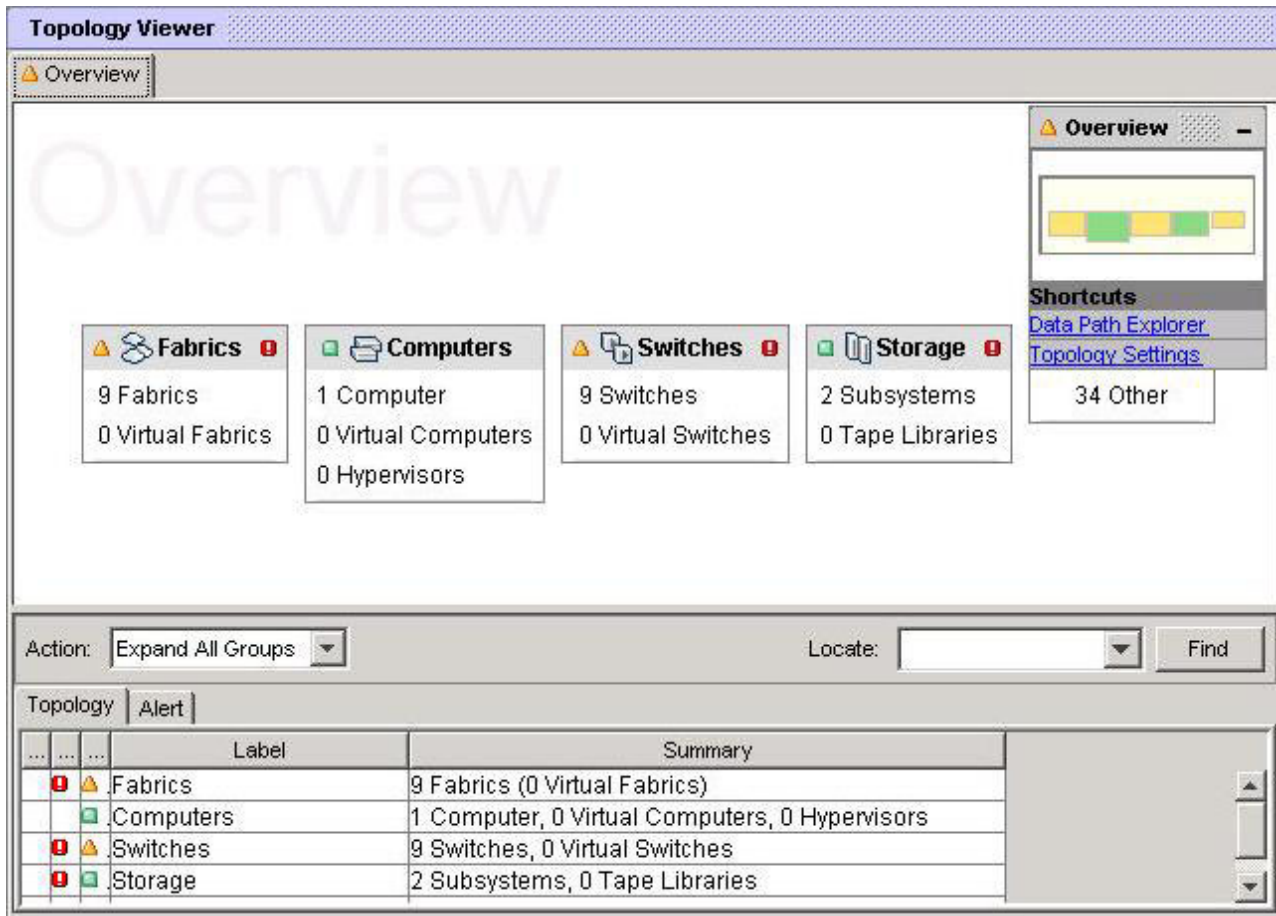


Figure 37. Topology overview

In addition, the topology viewer depicts the relationships among resources (for example, the disks comprising a particular storage subsystem). Detailed, tabular information (for example, attributes of a disk) is also provided. With all the information that the topology viewer provides, you can more quickly and easily monitor and troubleshoot your storage environment and perform critical storage-management tasks.

### Topology viewer interface controls

Progressive information disclosure reduces the visual complexity of a system. In the topology viewer, this means that you can identify and focus on entities in your storage environment that are of particular interest and hide other entities from view.

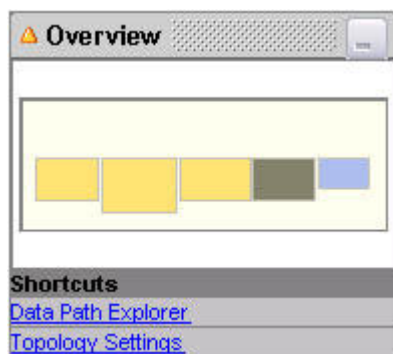
The tabs displayed in the tabular view represent classes of entities (a physical device or logical resource discovered by IBM Tivoli Storage Productivity Center). Individual entities (for example, a computer named *accounting1*, a tape library named *tapelib4*) are listed in their corresponding tabs. Entity attributes (for example, the operating system of a computer) are organized into table columns. The classes entities are assigned to are Computers, Fabrics, Storage Subsystems, and Other.

**Note:** The Other class contains entities discovered by Tivoli Storage Productivity Center but for which the type cannot be determined.

Changes to the graphical view or the tabular view are automatically applied to the other view, keeping both views synchronized. For example, if an entity is selected in the graphical view, it is also highlighted in the table view.

As more of your resources are discovered and displayed in the topology viewer, it is not possible to see all the discovered resources in a single topology window. To make it easier for you to find the resources you want to view, Tivoli Storage Productivity Center provides a tool called the *minimap*. The minimap allows you to shift the focus of the topology viewer to different areas within a particular view. The minimap is a small window that is initially located in the upper-right corner of any topology view.

**Linux/AIX users:** you must use SHIFT+ALT+mouse drag to reposition the minimap.



*Figure 38. Minimap.* You can click the shortcuts to launch Data Path Explorer or to view or change topology settings.

The minimap provides a miniature representation of the current topology view. Clicking on the minimap displays a multidirectional cursor that allows you to move the focus of the current topology view in any direction. As you drag the cursor around the minimap, the topology view will move accordingly. This makes navigating around complex topologies much easier.

The minimap uses the following colors to indicate the aggregated health status of groups:

- Green (normal): All entities in the group are operating normally.
- Yellow (warning): At least one entity is operating normally, and one or more entities have a health status of warning, critical, unknown, or missing.
- Red (critical): The status of all entities is either critical or warning.
- Gray (unknown): None of the entities in the group has an known health status.
- Black (missing): All of the entities of the group are missing.

In the Configuration History page, the minimap uses the following colors to indicate the aggregated change status of groups:

- Blue (changed): One or more entities in the group changed (an addition or deletion of an entity is considered a change).
- Gray (unchanged): All entities in the group are unchanged.

Included at the bottom of the minimap are shortcuts that you can click to launch the Data Path Explorer or to view or change the topology settings:

- To launch the Data Path Explorer, select one or more hosts, subsystems, disks, volumes, or Mdisks and click the Data Path Explorer shortcut.

**Note:** The shortcut to Data Path Explorer is not available in the Configuration History topology view.

- To view or change the topology settings, click the Topology Settings shortcut. The Settings window displays for you to specify the type of overlays to display, the type of filtering to perform, whether connections between topology objects should display, and whether the Zone tab should display in the table view.

### The graphical view:

This topic describes how the topology viewer groups entities that share certain characteristics. A collapsed group consisting of many entities or subgroups takes up much less space in the graphical view than the individual entities or subgroups would.

In the topology viewer, you can expand or collapse a group or subgroup by clicking the plus (+) character or minus (-) character, respectively, located in the upper-right corner of a group box.

**Note:** Pinned entities within a group remain visible, even when the group is collapsed.

By default, the topology viewer groups entities by class (computers, fabrics, switches, storage subsystems, and other entities). However, you can create additional groups (but not additional classes) using your own criteria. These custom groups can provide new perspectives on and important information about your storage environment. For example, you can create a special group belonging to the computer class that consists of computers that use the same operating system or storage subsystems of a particular model.

The following figure shows an expanded computer group representation in the graphical view:




Figure 39. Computer group

The title bar of the group box shows, from left to right, the following:

- An icon indicating the health status of the group (for example, normal).
- An icon indicating the types of entities contained in the group.
- The class of group (for example, computers).
- The grouping criterion (for example, Windows).
- The total number of entities in the group. If filtering is turned on, the number in the brackets is the ratio of displayed entities to total entities. For example, 5/10 indicates that 5 entities are shown out of 10 total entities.

The topology viewer provides the following default groups. These groups are accessible, if they exist in your topology, in the overview or by expanding the primary default groups (**Computers**, **Fabrics**, **Switches**, **Storage** and **Other**) that are displayed when you first launch the topology viewer.

Table 25. Group locations in the topology viewer

| Type of group  | Description                                                                                                                      | Path to group in topology viewer                                                                                                                               |
|----------------|----------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Computers      | All the computers that are discovered in your environment or a group of computers associated with a particular fabric or switch  | <b>Computers</b>                                                                                                                                               |
| Fabrics        | All the fabrics and virtual fabrics that are discovered in your environment                                                      | <b>Fabrics</b>                                                                                                                                                 |
| Storage        | All the storage-related devices that are discovered in your environment                                                          | <b>Storage</b>                                                                                                                                                 |
| Other          | All the entities that are discovered in an environment but that do not fit in either the computer, fabric, or storage categories | <b>Other</b>                                                                                                                                                   |
| Fabric         | All entities (for example, switches) that belong to or are related to a fabric                                                   | <b>Fabrics &gt; Fabric</b>                                                                                                                                     |
| Virtual Fabric | All entities that belong to a virtual fabric                                                                                     | <b>Fabrics &gt; Fabric</b>                                                                                                                                     |
| Switches       | A set of switches                                                                                                                | <b>Fabrics &gt; Fabric</b>                                                                                                                                     |
| Subsystems     | A set of subsystems                                                                                                              | <b>Storage &gt; Subsystems</b>                                                                                                                                 |
| Tape libraries | A set of tape libraries                                                                                                          | <b>Storage &gt; Tape Libraries</b>                                                                                                                             |
| Device         | Devices associated with a particular entity (for example, a computer)                                                            | <b>Computers &gt; Computers</b>                                                                                                                                |
| Volumes        | A set of volumes that are associated with an entity                                                                              | Different types of volumes groups exist at the L2 level of each of the four main categories: <b>Computers</b> , <b>Fabrics</b> , <b>Storage</b> , <b>Other</b> |
| HBA            | A set of HBAs that are associated with an entity                                                                                 | <b>Computers &gt; Computers &gt; Computer</b>                                                                                                                  |
| Disks          | A set of disks that are associated with an entity                                                                                | <b>Storage &gt; Subsystems &gt; Subsystem</b>                                                                                                                  |
| Pools          | A set of pools that are associated with an entity                                                                                | <b>Storage &gt; Subsystems &gt; Subsystem</b>                                                                                                                  |
| FC Ports       | A set of FC ports that belong to an entity                                                                                       | <b>Fabrics &gt; Fabric &gt; Switch</b>                                                                                                                         |
| Media          | All media attached to an entity                                                                                                  | <b>Storage &gt; Tape Libraries &gt; Tape Library</b>                                                                                                           |
| Media Changers | All media changers attached to an entity                                                                                         | <b>Storage &gt; Tape Libraries &gt; Tape Library</b>                                                                                                           |
| Tape Drives    | All tape drives attached to an entity                                                                                            | <b>Storage &gt; Tape Libraries &gt; Tape Library</b>                                                                                                           |

Table 25. Group locations in the topology viewer (continued)

| Type of group | Description                                                                                                                   | Path to group in topology viewer        |
|---------------|-------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------|
| Tape Drive    | All ports on a tape drive attached to an entity                                                                               | Storage > Tape Libraries > Tape Library |
| Switch        | All ports on a switch. Due to the large number of ports on some switches, ports are shown as FC port groups within the switch | Fabrics > Fabric > Switch               |

As shown by the tabs at the top of the graphical panel (**L0:Computers**, **L1:Computers**, and **L2:Computers**), other views have been opened but are hidden by the overview. The topology tab in the table view displays information about the entities in the groups. The minimap, which you can use to shift the focus in the view, is positioned in the upper-right corner of the graphical view.

Right-click the object and select **Open Detail View** to view entities in a more detailed view. For example, at a high level an object representing a storage subsystem is depicted as an icon with a label. At lower levels, the same storage subsystem can be rendered as a box containing sub-entities representing disks, with the corresponding table view showing attributes of each individual disk. At yet another detail level, the ports of the storage subsystem might be shown as well as status information for each port.

*GUI Topology Viewer has new L0 and L1 Switches views:*

The IBM Tivoli Storage Productivity Center GUI has a Switches node under the Topology node. The Switches node opens the new L0 Switches view, from which you can also access the new L1 Switches view.

### Overview

The Overview panel now shows a box with the number of switches present in your environment.

### L0 Switches

The L0 Switches view shows all switches in your environment.

If you right-click a Switches group, then click Group By, you will get a dialog box for group settings. You can choose to group by a single group or you can group by multiple categories:

- ByHealth
- ByUDP1
- ByUDP2
- ByUDP3
- ByFabric
- ByGroupSize

Click OK to close the dialog box and update the group settings.

### L1 Switches

The L1 Switches view shows information on a group of switches in your environment that you selected to display. This also includes connectivity information: what entities this group is connected to.

### L2 Switch

This view shows one switch, with all detailed port and connectivity

information. You can access this view by double-clicking a single switch icon in either the L0 or L1 Switches view, or through the other Topology Viewer views.

**The table view:**

This topic describes how tabs in the table view provide detailed information about the entities (for example, computers, subsystems and so on) that are relevant in the current view.

The following is a list of all the tabs available in the topology viewer. Not all tabs are visible in all views.

**Topology**

Detailed information about all the entities in the topology grouped by class of the entity (computers, fabric, storage, and other).

**Computer**

Detailed information about any computer in the computer group or groups in the current topology view.

**Fabric** Detailed information about a fabric groups and virtual fabric groups in the current topology view.

**Switch**

Detailed information about switch groups and switch entities in the current topology view.

**Subsystem**

Detailed information about volume groups and volume entities in the current topology view.

**Disk** Detailed information about disk groups and disk entities in the current topology view.

**Pool** Detailed information about pool groups and pool entities in the current topology view.

**Note:** One of the columns in the tabular view refers to whether or not a storage pool is primordial. A primordial pool represents unallocated storage capacity on a storage device. Storage capacity can be allocated from primordial pools to create storage pools. This means that primordial pools are disk/device sources for allocation of storage pools.

**Volume**

Detailed information about volume groups and volume entities in the current topology view.

**Tape Library**

Detailed information about tape-library groups and tape-library entities in the current topology view.

**Media Changer**

Detailed information about media-changer groups and media-changer entities in the current topology view.

**Tape Drive**

Detailed information about tape-drive groups and tape-drive entities in the current topology view.

**Other/Unknown**

Detailed information about objects in your environment other than computers, switches, subsystems, and tape libraries in the current topology view.

**FC Port**

Detailed information about port groups and port entities in the current topology view.

**Zone** Detailed information about any of the following groups or entities in a zone. By default, the **Zone** tab is disabled.

- Fabric groups
- Zone-set groups
- Zone groups
- Computer entities
- Switch entities
- Subsystem entities
- Tape library entities













**Alert** Detailed information about all alerts that are associated with entities in the current view of the topology.

**Identifying storage resources using icons:**

This topic describes how the topology viewer displays entities and groups of entities with an appropriate icon based on the entity class.

The storage resource icons can help you quickly and easily identify storage resources. For example, an entity of class "computer" is displayed in the graphical view with an icon that resembles a computer.

The following table shows the icons for individual resources and for resource groups.

| Topology Object   | Icon                                                                                | Group Icon                                                                            |
|-------------------|-------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| Computer          |  |  |
| Computer HBA      |  |  |
| Fabric            |  |  |
| Switch            |  |  |
| Switch port       |  |  |
| Storage subsystem |  |  |

| Topology Object                                                          | Icon | Group Icon |
|--------------------------------------------------------------------------|------|------------|
| Storage subsystem port                                                   |      |            |
| Storage virtualizer (for example, SAN Volume Controller, HDS Tagmastore) |      |            |
| Tape library                                                             |      |            |
| Tape library media changer                                               |      |            |
| Database volume                                                          |      |            |
| File system volume                                                       |      |            |
| Storage volume                                                           |      |            |
| Other volume                                                             |      |            |
| Storage disk                                                             |      |            |
| Storage backend disk                                                     |      | n/a        |
| Storage pool                                                             |      |            |
| Unknown                                                                  |      |            |







**Note:** *Unknown* entities are entities that are discovered by IBM Tivoli Storage Productivity Center but that do not have an agent installed or the agent is not working.

#### Identifying icons for storage resources that contain full disk encryption disks

The icons in the following table identify whether a storage system contains encrypted or encryptable disks. The group icon for encrypted storage disks, pools, and volumes is the same as the group icons that are shown in the first table.

For more information on supported systems that provide full disk encryption, see “Supported storage systems providing full disk encryption and solid-state drives” on page 1145.





| Topology Object | Encrypted Icon                                                                    | Encryptable Icon                                                                    |
|-----------------|-----------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|
| Storage disk    |  |  |
| Storage pool    |  |  |
| Storage volume  |  |  |

### Identifying icons for storage resources that contain solid state drives

A solid-state drive is a storage device that stores persistent data on solid-state flash memory. The icons in the following table identify whether a storage system contains solid-state drives. The group icons for storage disk, pool, and volume in the following table are the same as the group icons that are shown in the first table.

For more information on supported systems that provide solid-state drives, see “Supported storage systems providing full disk encryption and solid-state drives” on page 1145.

| Topology Object                | Icon                                                                                |
|--------------------------------|-------------------------------------------------------------------------------------|
| Solid-state drive storage disk |   |
| Solid-state drive storage pool |  |

### The topology viewer and virtual SANs:

This topic describes how Cisco switches support virtual SANs (VSANs), which is the logical partitioning of a fabric into multiple fabrics. The overall network is referred to as the physical infrastructure, and the logical fabrics are the VSANs.

SNMP queries to a Cisco switch collect information about the Cisco physical infrastructure and about any VSANs in that infrastructure. Storage Resource agents or In-band fabric agents can also collect information about the VSANs but are unable to collect information about the physical infrastructure. The status of Cisco physical infrastructures shown in the topology view is determined based on the status propagated from the VSANs in the infrastructure and the status of the infrastructure itself. If the status of the physical infrastructure is missing but the status of all or some of the VSANs in the infrastructure is normal, infrastructure information can no longer be collected. This typically indicates that the out-of-band agent is no longer working properly or that SNMP queries to the Cisco switches are no longer configured because the switch has been removed from the list of agents in the **Out of Band Fabric** node.

The consequences of this situation are illustrated in the following scenario: A fabric includes a Cisco switch with a VSAN on it. An out-of-band agent on the switch collects physical infrastructure and VSAN information, and an in-band Fabric agent collects information about the VSAN. If the out-of-band agent is removed or

Fabric Manager is no longer able to communicate with the agent, IBM Tivoli Storage Productivity Center still detects the VSAN through Storage Resource agents or in-band agents, but physical infrastructure information can no longer be collected. As result, in the fabric topology view, the status of the physical infrastructure box is shown as missing, while the status of one or more VSANs in the infrastructure shows as normal. To correct this situation ensure that the out-of-band agent is configured for the Cisco switch.

### **The topology viewer and VMware:**

This topic describes what the topology viewer displays for VMware support.

The Overview view of the topology viewer displays a count of hypervisors and virtual machines in the Computers summary.

The L0: Computers view displays the hypervisors and virtual machines in addition to other computers and includes the grouping By Virtualizer.

The L1: Computer view for a hypervisor displays volumes that were assigned to the hypervisor and the connectivity to fabrics.

The L1: Computer view for a virtual machine displays volumes that were assigned to the hypervisor and have been mapped to the VMware computer (virtual machine) as raw disks.

The L2: Computer view for a hypervisor includes a mapping of virtual machines and shows the virtual machine disk to physical disk relationship.

The L2: Computer view for a virtual machine displays the physical disk (what the virtual machine recognizes as a physical disk) to virtual machine disk relationship and hypervisor connectivity.

**Note:** IBM Tivoli Storage Productivity Center supports the mapping of storage between the storage system volume and ESX Server corresponding disk. This mapping is supported for the ESX Server 3.5 or VirtualCenter 2.5 or later. Previous versions of ESX Server or VirtualCenter are not supported for the mapping of storage to the storage system.

### **The topology viewer and switches:**

Learn about important considerations when viewing and collecting information about McDATA and Brocade switches in the topology viewer.

#### **McDATA Switches, CIMOMS, and their status as displayed through the topology viewer**

McDATA CIMOMS communicate with all switches in a fabric as long as they are all managed by the Enterprise Fabric Connectivity Management (EFCM) Server and the CIMOM is pointed at the EFCM Server. Alternatively, if a McDATA CIMOM is configured in "Direct Mode", it communicates with all switches in a fabric as long as the CIMOM has been configured to manage all switches.

When a switch is disabled, the CIMOM can still communicate with the switch so it does not go missing. McDATA leaves the operational status and health indications the same for the switch and then changes the operational status for all its ports to dormant.

Therefore, only the port-to-port connections appear in the topology viewer as missing, as well as any nodes connected to those ports. The ports might have a yellow status indicator due to their dormant status. This occurs whether the fabric is set up in a flat or a loop/mesh topology.

### **Brocade Switches, CIMOMS, and their status as displayed through the topology viewer**

Brocade CIMOMs do not communicate directly with all switches in a fabric even when those switches are all managed by the CIMOMs. Brocade fabrics use a "proxy switch" mechanism for communication. This proxy switch is the only switch in the fabric to which the CIMOM communicates.

When the proxy switch is disabled, the CIMOM can still communicate with the switch so it does not go missing. Brocade leaves the operational status and health indications the same for the switch and all its ports, but then changes the **Enabled State** for all its ports to **disabled**. Therefore, the switch, its ports, and all its port-to-port connections appear in the topology viewer with the same **Operational Status** as they had before the switch was disabled.

However, because the proxy switch is disabled, the rest of the switches in the fabric go missing, as well as any nodes connected to switches in the fabric. This occurs whether the fabric has its switches in a flat or a loop/mesh topology.

When a "non-proxy switch" is disabled, the CIMOM can no longer communicate with that switch using the proxy switch. Therefore, the switch, its ports, and all its port-to-pot connections appear in the topology viewer as missing. For a Brocade fabric with a loop/mesh topology only, the disabled switch is affected because the proxy switch can reach the other switches through other ISL connections.

To workaroud this issue, it is recommended you set up redundant agents (such as Out of Band Fabric Agents) for the Brocade switches in your fabrics.

### **Viewing large environments in the Topology Viewer:**

In environments where you are monitoring a large number of storage entities, IBM Tivoli Storage Productivity Center might not be able to display all of those entities within the Topology Viewer. To avoid this issue, there are some tasks you can perform to improve overall scalability and performance of the user interface.

If the Topology Viewer is unable to display all the entities in your environment, the user interface generates javacore and heapdump files in its installation directory. It is recommended that you check this directory to determine if this display problem is occurring in your environment. To do this, perform the following steps:

1. View the contents of the installation directory where Tivoli Storage Productivity Center's user interface is installed. The default installation directories are as follows:
  - Windows: C:\Program Files\IBM\TPC\gui\
  - UNIX/Linux: /<usr or opt>/IBM/TPC/gui/
2. Determine if javacore and heapdump files appear in the installation directory. You can identify these files by searching for files that begin with the word **javacore** or **heapdump** and have a **.txt** or **.phd** extension. For example: javacore.20080605.151906.8088.txt and heapdump.20080605.151903.8088.phd.

If these files appear in the installation directory, it might indicate that the display problem is occurring when you use the Topology Viewer.

There are two methods for working around this issue:

- Reducing the number of open topology views and disable unneeded overlays (such as the alert, performance, and zoning overlays). The number of topology views that are open and overlays that are enabled can affect how many storage entities are able to be displayed in the Topology Viewer.
- Increasing the memory allocation for IBM Tivoli Storage Productivity Center's user interface component. If reducing the number of open topology views and disabling unneeded overlays does not correct the problem, you can increase the amount of memory that is allocated to the user interface to help resolve the issue.

*Reducing the number of open topology views and disabling unneeded overlays:*

Perform these tasks to help improve the scalability of the Topology Viewer.

The number of topology views that are open can affect how many storage entities are able to be displayed within those views. The more views (tabs) that are open in the Topology Viewer, the more memory is used by the user interface. If the user interface requires more memory than is available to display all the entities within the open views of the Topology Viewer, some of the entities are not displayed. To work around this problem, close any views that you no longer need to view and refresh the content of the views that you leave open.

The number of overlays that are enabled in the Topology Viewer can also affect how many storage entities are able to be displayed. The topology viewer uses color and icons to indicate the status of objects in the topology and tabular views. Overlays are colors and icons that used to indicate health, performance, change, and alerts. To disable overlays that you do not need, complete the following steps:

1. Right-click any blank space within a topology view.
2. Select **Global Settings...** from the pop-up menu.
3. Uncheck the overlays you do not need to view. For example, clear the **Performance** and **Alert** check boxes.
4. Click **OK** to accept the changes. The content of the open views within the Topology Viewer are automatically refreshed according to your changes.

*Increasing the memory allocation for the user interface:*

Perform this task to allocate more system memory to the IBM Tivoli Storage Productivity Center user interface.

Increasing the amount of allocated memory might enable the user interface to correctly display a large number of storage entities in the Topology Viewer. This task is generally recommended and encouraged in larger environments.

To do this, perform the following steps:

1. Close any current sessions of the Tivoli Storage Productivity Center user interface.
2. Locate the batch or script file that is used to start the user interface. This file is located in the Tivoli Storage Productivity Center installation directory on the computer where the user interface is installed. The default file names and installation directories are as follows:

- Windows: C:\Program Files\IBM\TPC\gui\tpc.bat
  - UNIX/Linux: /<usr or opt>/IBM/TPC/gui/TPCD.sh
3. Open the tpc.bat or TPCD.sh file in a text editor and locate both occurrences of the parameter that defines how much memory is allocated to the GUI component:
    - Windows parameter: **-mx256M**. The default setting for Windows is 256 megabytes.
    - UNIX/Linux parameter: **-xmx512M**. The default setting for UNIX and Linux is 512 megabytes.
  4. Increase the value for all occurrences of this parameter. For example:
    - Windows: change **-mx256M** to **-mx1024M**
    - UNIX/LINUX: change **-xmx512M** to **-xmx1024M**

**Note:** Make sure that the computer on which the user interface is installed has enough memory to support the setting that you define.

5. Save the batch or script file and restart the user interface.
6. Check the Topology Viewer to ensure that all the monitored storage entities are displayed as a result of increasing the amount of memory allocated to the user interface. If you still experience this display problem, repeat steps 1-5 and further increase the value for the parameter in the startup file.

Contact an IBM service representative if changing the memory allocation for the user interface does not help resolve the issue.

## Arranging entities in a group

This topic describes how to arrange entities in an expanded group.

You can arrange entities within the group box by health status or by entity label if an expanded group contains many entities. With either type of arrangement, you can specify ascending or descending order. To arrange entities, follow this procedure:

1. Right-click the graphical view and select **Arrange By**.
2. Select one of the following:
  - **Label (ascending)**
  - **Label (descending)**
  - **Health (ascending)**
  - **Health (descending)**

## Checking the status of entities and groups

This topic describes how the topology viewer uses icons to indicate the health, performance, and operational status for entities.

To quickly check the health, performance, operational, change, and alert status of entities or groups of entities, look for the following icons associated with a particular entity or group in the topology. Position the cursor over these icons to see the health string in text format:

*Table 26. Topology status indicators and their icons*








| Indicator       | Icon                                                                                  |
|-----------------|---------------------------------------------------------------------------------------|
| Health - Normal |  |

Table 26. Topology status indicators and their icons (continued)

| Indicator                                                              | Icon                                                                                  |
|------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| Health - Warning                                                       |    |
| Health - Critical                                                      |    |
| Health - Unknown                                                       |    |
| Health - Missing                                                       |    |
| Operational - Running                                                  |    |
| Operational - Upgrading                                                |    |
| Operational - Missing                                                  |    |
| Performance - Normal                                                   |    |
| Performance - Unknown                                                  |    |
| Performance - Warning                                                  |    |
| Performance - Critical                                                 |   |
| Performance - Missing                                                  |  |
| Change - Objects did not exist at the two specified time periods       | Objects in light gray with no icon                                                    |
| Change - Objects changed between the two specified time periods        |  |
| Change - Objects did not change between the two specified time periods |  |
| Change - Objects were created between the two specified time periods   |  |
| Change - Objects were deleted between the two specified time periods   |  |
| Alert - Critical, Warning, or Informational                            |  |

The operational status icons (running, upgrading, missing) are displayed by default and cannot be hidden. However, you can specify whether to display health, performance, or alert icons. To display these icons, follow this procedure:

1. Right-click a blank space in the background of the graphical panel and select **Global Settings**.
2. In the Settings window, select **Health**, **Performance**, or **Alert**. Click **OK**.

### Health status

This topic describes how the topology viewer displays the health status of entities.

The health of an entity falls into one of the following categories:

- Normal: The entity is operating normally.
- Warning: At least one part of the entity is not operating or has serious problems.
- Critical: The entity is either not operating or has serious problems.
- Unknown: The entity was discovered but is not recognized by IBM Tivoli Storage Productivity Center.
- Missing: The entity was previously discovered by Tivoli Storage Productivity Center, but was not detected in the most recent discovery.

The health of a group falls into one of the following four categories:

- Normal: All the entities in the group are operating normally.
- Warning: At least one entity in the group has normal status.
- Critical: All entities are either in critical or warning states.
- Unknown: None of the entities in the group have a known health status.

Health indicators for some entities, such as ports, are only normal and critical. In certain situations, the health of an entity might not be available. This happens, for example, when the entity is identified as missing. In the graphical view, missing entities are shown with a semitransparent light gray layer on top of an existing icon.

Level-indicator buttons at the top of the graphical view (for example, **L0:Computers**) display icons for health.

Multiple states from single entities are aggregated into group states or states of entities shown in upper layers, using the following rules:

- If all entities are in the same state, the aggregated state will represent this state.
- You will get an aggregated critical state (red) if:
  - One entity is in the missing state and at least one other entity shows a critical state.
  - One entity is in the critical state and at least one other entity shows a missing state.
- You will get an aggregated normal state if:
  - One entity is in the normal state and all others are in normal or undefined state.
  - One entity is in the undefined state and all others are in normal or undefined state.
- All other combinations will aggregate to warning state (yellow).
- States displayed within the tabs on top of the topology view will be aggregated according to the focus of the specific entity class.
  - For Computers and Storage, only the states of these entities are used for aggregation.
  - For Fabric, only the switch status is used for aggregation.
  - For level 1 and 0, the states of all displayed entities are used for aggregation.

See “How health status is determined for individual storage entities” on page 530 for more information about how the status of individual entities are displayed in the topology viewer.

At the Overview level, the health status indicated by the icon on **Overview** at the top of the graphical view indicates the aggregated status of all the groups inside the graphical view (**Computers, Fabrics, Switches, Storage and Other**). At the L0

level, the health status indicated by the icon on the **L0** tab at the top of the graphical view indicates the aggregated status of all the groups inside the that particular L0 graphical view (for example, the health status for the L0:Computers view would aggregate the health status of all discovered computers, including those that are missing, computers whose health status is critical, computers whose health status is warning). At the L1 level, the health status indicated by the icon on the **L1** tab at the top of the graphical view indicates the status of the particular group that is the focus of the view. (The status of the group is based on the aggregated status of the entities within the group.) The health status for the group is also indicated in the table view. At the L2 level, the health status indicated by the icon on the **L2** tab at the top of the graphical view is the same as the status of the particular entity that is the focus of the view. The health status for the entity is also indicated in the table view.

**Note:**






- At the L1:Fabric level the health status indicated by the icon on the **L1:Fabric** tab at the top of the graphical view might not be the same as the overall health of the group, based upon the aggregated status of all the entities within the group.
- At the L2:Subsystem level, the status of the volume group is the aggregated status of all the volumes that comprise the group and not the status of the host that owns those volumes.
- The number of computers that are displayed in the Computers view of the topology viewer might be more than the number of Tivoli Storage Productivity Center agents that are deployed in the environment. These additional nodes are displayed with an "unknown" status. For example, if Tivoli Storage Productivity Center determines that a node in a fabric is an HBA, that node is displayed in the Computers view (rather than in the Other view). This situation occurs because during a probe, inband fabric agents can identify nodes in the topology viewer as HBAs in computers.

To update nodes that have an unknown status in the Computers view, deploy an Storage Resource agent on the computer that has the HBA whose WWN is shown in the topology viewer.

**How health status is determined for individual storage entities:**

The topology viewer displays the health status for individual storage entities and for groups of storage entities. Use this topic to learn how Tivoli Storage Productivity Center determines the health statuses for individual storage entities.

The following health statuses are possible for an individual storage entity:

-  Normal: The storage entity is operating normally.
-  Warning: At least one part of the storage entity is not operating or has serious problems.
-  Critical: The storage entity is either not operating or has serious problems.
-  Unknown: The storage entity was discovered but is not recognized by Tivoli Storage Productivity Center.
-  Missing: The storage entity was previously discovered by Tivoli Storage Productivity Center, but was not detected in the most recent discovery.

Tivoli Storage Productivity Center determines the health status of storage entity by propagating the health status of its related components. The status of the following components determines the propagated status of the related storage entity:



### **Storage subsystems**

The health statuses of the ports, physical disks, pools, volumes, switches, and consolidated status related to a storage subsystem are used to calculate the overall health status of that storage subsystem in the topology viewer. For example, if the status of some volumes is critical and the status of disks is unknown, then the health status of the storage subsystem is displayed as critical.

For Storwize V7000 File Module, Network Shared Disk (NSD), node, and cluster status are also included to calculate the overall health status of the storage system in the topology viewer.

### **Computers**

The health statuses of the HBAs, controller ports, and consolidated status for a computer are used to calculate the overall health status of that computer in the topology viewer. For example, if the status of an HBA is unknown and the status of the controllers are normal, then the health status of the computer is displayed as normal.

**HBAs** The health statuses of the ports and consolidated status for an HBA are used to calculate the overall health status of that HBA in the topology viewer.

### **Switches**

The health statuses of the ports and consolidated status for a switch are used to calculate the overall health status of that switch in the topology viewer.

### **Fabrics**

The following health statuses are used to calculate the overall health status that is displayed for a fabric in the topology viewer:

- The health of the “other” entities part of the fabric
- The health of the computer entities part of the fabric
- The health of the storage subsystem entities part of the fabric
- The health of the tape library entities that are part of the fabric
- The health of the switch entities that are part of the fabric
- The consolidated status of the fabric

### **Tape drives**

The health statuses of the node entities that are connected to a tape drive and the consolidated status of a tape drive are used to calculate the overall health status of a tape drive in the topology viewer.

### **Tape media changers**

The health statuses of the node entities that are connected to a tape media changer and the consolidated status of a tape media changer are used to calculate the overall health status of a tape media changer in the topology viewer.

### **Tape libraries**

The following health statuses are used to calculate the overall health status that is displayed for a tape library in the topology viewer:

- The health statuses of the node entities that are connected to the tape library
- The health of all the tape drive entities connected to the tape library
- The health of all the tape media changer entities connected to the library
- The consolidated status of the tape drive

Use the following table to learn how Tivoli Storage Productivity Center uses the health status of components to calculate the health status for the related storage entity. The **Unknown**, **Normal**, **Critical**, **Warning**, and **Missing** columns indicate the possible statuses of components, and the value in the **Resulting health status** column represents the propagated status for the related storage entity.

Table 27. Calculation of health status for a storage entity










































| Unknown  | Normal  | Critical  | Warning  | Missing  | Resulting health status                                                                        |
|-------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------|
| X                                                                                         |                                                                                          |                                                                                            |                                                                                            |                                                                                             | Unknown     |
| X                                                                                         | X                                                                                        |                                                                                            |                                                                                            |                                                                                             | Normal      |
| X                                                                                         | X                                                                                        | X                                                                                          |                                                                                            |                                                                                             | Warning     |
| X                                                                                         | X                                                                                        | X                                                                                          | X                                                                                          |                                                                                             | Warning     |
| X                                                                                         | X                                                                                        | X                                                                                          | X                                                                                          | X                                                                                           | Warning     |
| X                                                                                         | X                                                                                        | X                                                                                          |                                                                                            | X                                                                                           | Warning     |
| X                                                                                         | X                                                                                        |                                                                                            | X                                                                                          |                                                                                             | Warning     |
| X                                                                                         | X                                                                                        |                                                                                            | X                                                                                          | X                                                                                           | Warning     |
| X                                                                                         | X                                                                                        |                                                                                            |                                                                                            | X                                                                                           | Warning     |
| X                                                                                         |                                                                                          | X                                                                                          |                                                                                            |                                                                                             | Warning     |
| X                                                                                         |                                                                                          | X                                                                                          | X                                                                                          |                                                                                             | Warning     |
| X                                                                                         |                                                                                          | X                                                                                          |                                                                                            | X                                                                                           | Warning     |
| X                                                                                         |                                                                                          | X                                                                                          | X                                                                                          | X                                                                                           | Warning    |
| X                                                                                         |                                                                                          |                                                                                            | X                                                                                          |                                                                                             | Warning   |
| X                                                                                         |                                                                                          |                                                                                            |                                                                                            | X                                                                                           | Warning   |
| X                                                                                         |                                                                                          |                                                                                            | X                                                                                          | X                                                                                           | Warning   |
|                                                                                           | X                                                                                        |                                                                                            |                                                                                            |                                                                                             | Normal    |
|                                                                                           | X                                                                                        | X                                                                                          |                                                                                            |                                                                                             | Warning   |
|                                                                                           | X                                                                                        | X                                                                                          | X                                                                                          |                                                                                             | Warning   |
|                                                                                           | X                                                                                        | X                                                                                          |                                                                                            | X                                                                                           | Warning   |
|                                                                                           | X                                                                                        | X                                                                                          | X                                                                                          | X                                                                                           | Warning   |
|                                                                                           | X                                                                                        |                                                                                            | X                                                                                          |                                                                                             | Warning   |
|                                                                                           | X                                                                                        |                                                                                            | X                                                                                          | X                                                                                           | Warning   |
|                                                                                           | X                                                                                        |                                                                                            |                                                                                            | X                                                                                           | Warning   |
|                                                                                           |                                                                                          | X                                                                                          |                                                                                            |                                                                                             | Critical  |
|                                                                                           |                                                                                          | X                                                                                          | X                                                                                          |                                                                                             | Warning   |

Table 27. Calculation of health status for a storage entity (continued)

| Unknown  | Normal  | Critical  | Warning  | Missing  | Resulting health status                                                                      |
|-------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------|
|                                                                                           |                                                                                          | X                                                                                          | X                                                                                           | X                                                                                           | Warning   |
|                                                                                           |                                                                                          | X                                                                                          |                                                                                             | X                                                                                           | Critical  |
|                                                                                           |                                                                                          |                                                                                            | X                                                                                           |                                                                                             | Warning   |
|                                                                                           |                                                                                          |                                                                                            | X                                                                                           | X                                                                                           | Warning   |
|                                                                                           |                                                                                          |                                                                                            |                                                                                             | X                                                                                           | Missing   |

## Operational status

This topic describes how the topology viewer has three categories of operational status for entities.

The operational status of most entities falls into one of the following three categories:

Table 28. Categories of operational status for storage entities

| Status    | Description                                                  |
|-----------|--------------------------------------------------------------|
| Running   | The entity is running.                                       |
| Upgrading | The entity is upgrading.                                     |
| Missing   | The entity was not discovered during the last refresh cycle. |

However, some entities display additional operational status. For example, switches can display a status of **OK**, **Error**, **Stopped**, and other types of status.

**Note:** When IBM Tivoli Storage Productivity Center no longer detects an entity that was previously displayed in the Topology viewer, the operational status of that entity remains as the last known status before it went missing. This means that the operational status of an entity is not cleared or reset if that entity is not found in subsequent data collection jobs.

## Performance status

This topic describes performance status details collected for entities.

Performance information is collected for entities in the L2:Switch and L2:Subsystem views. You can see the performance indicators for individual switches and subsystems in these two views by turning on the performance feature in the Global Settings window. This feature is not enabled by default.

The performance status of an entity falls into one of the following five categories:

- Normal: The entity is performing at the expected level.
- Warning: At least one part of the entity is not operating at the expected performance level.
- Critical: The entity is operating below the expected performance level.
- Unknown: The performance of the entity is not known.
- Missing: The entity was not found during the last refresh cycle and the performance status is not available.

The performance status does not get carried up to the L0 and L1 views. It is only available in the L2:Switch and L2:Subsystem views.

Performance status of an entity is not propagated upwards to its parent entity. For example, a SAN Volume Controller can cross several critical or warning stress thresholds (and show as yellow or red in the topology viewer) even though the underlying pools, disks or volumes do not have a degraded performance status.

After performance has been enabled, you will see a performance status indicator for any entity in these views that has performance data available. Position the cursor over a performance icon to see a list of performance metrics for that entity.





To enable the collection of performance information, perform the following steps:

1. Right-click a blank space in the background of the graphical panel and select **Global Settings**.
2. In the Settings window, select **Performance** and click **OK**.

## Change status

This topic describes how the Configuration History topology view displays the status of entities over time.

The Configuration History feature uses the icons and colors of the change overlay to indicate historical changes in a SAN configuration's entities. The same icons and colors are used in the graphical view, table view, and snapshot selection panel. The following table lists the five possible states of an entity and shows the associated icons and colors.

| State of an entity                                                                                                       | Change overlay icon                                                                                        | Change overlay color indicator |
|--------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------|--------------------------------|
| Entity changed between the time that the snapshot was taken and the time that a later snapshot was taken.                | <br>Yellow pencil     | Blue background                |
| Entity did not change between the time that the snapshot was taken and the time that a later snapshot was taken.         | <br>Light gray circle | Dark gray background           |
| Entity was created or added between the time that the snapshot was taken and the time that a later snapshot was taken.   | <br>Green cross       | Green background               |
| Entity was deleted or removed between the time that the snapshot was taken and the time that a later snapshot was taken. | <br>Red minus sign    | Red background                 |
| Entity did not exist at the time that the snapshot was taken or at the time that a later snapshot was taken.             | Not applicable                                                                                             | Light gray background          |

Level-indicator buttons at the top of the graphical view (for example, **L0:Computers**) display the change icons.

For groups and other collections, the change overlay provides only aggregated status. That is, the change status for a group of entities is either “changed” or “unchanged.” Details for the group as a whole are not provided. Aggregated status displays in the minimap. In the Configuration History page, the minimap uses the following colors to indicate the aggregated change status of groups:

- Blue (changed): One or more entities in the group changed (an addition or deletion of an entity is considered a change).
- Gray (unchanged): All entities in the group are unchanged.

An object that is moved from one group to another group appears only once in the new group in the graphical view. For example, if the health status of a computer has changed from Normal to Warning, the Configuration History page displays the computer as changed in the Warning health group (and no longer displays the computer in the Normal health group).

**Important:** If you manually deleted a missing entity from the Tivoli Storage Productivity Center, the configuration history function rediscovers and displays the entity in the configuration history viewer as both Missing and in a Normal state. For example, Entity A (with ID 12) is missing and you decide to manually delete it by using the Remove From Database option from the pop-up menu for Entity A in the topology viewer or configuration history viewer. Later, the Tivoli Storage Productivity Center rediscovers Entity A and assigns it a new ID (ID 22). The configuration history function treats the same Entity A (with ID 12 and ID 22) as two different entities and displays them in both Missing and Normal states in the configuration history viewer. This scenario can also occur if a missing entity is removed by the Tivoli Storage Productivity Center removed resource retention function and is rediscovered by the Tivoli Storage Productivity Center again at a later time.

To display a summary of change for the entire row, hover over the changed row. The summary displays the title of the column. It also displays the old and new value for each cell or tab that changed. For example, if a group contains one new entity but everything else in the group is unchanged, the group as a whole displays as “changed.” If you click the group, the one new entity displays as “created” and all other entities display as “unchanged.”

The Change column in the table view also indicates changed entities.


## Alert status

This topic describes how the topology viewer displays alert status for all entities in the current view.

The alert status of an entity falls into one of the following categories:

- Warning: At least one part of the entity is not operating or has serious problems.
- Critical: The entity is either not operating or has serious problems.
- Informational: Related data about the entity is supplied to aid you in decision making.

When enabled, in the topology viewer the alert overlay is located at the right of entity icons, below the pinning indicator. The icon for the alert overlay is a red

exclamation mark  for the presence of alerts. In the title bar of groups and tabs, the alert icon is similarly located at the far right of the entity label. The alert overlay is disabled in the Configuration History page.

To find alerts that are associated with an entity in the table view, click the alert overlay indicator to go to the corresponding alert below the Alert tab. Or, right-click the entity to display a pop-up menu and click **Show Alerts in Tabular View** to go to the Alert tab and any corresponding alerts that are associated with that entity. The corresponding alerts are highlighted in purple. When you select an

alert in the Alert tab, the entities that are associated with that alert are highlighted in purple in the graphical view. Note that more than one entity may become highlighted. An alert icon beside a collapsed group indicates that alerts exist within the collapsed group.

The alert overlay provides rollup, similar to pinning. If alerts are inside a collapsed group, the number of alerts within that group are aggregated and indicated in the collapsed group. All alerts are aggregated into one count, independent of type or entity involved. Hover over the alert indicator for a tab or group to determine the total number of alerts that are rolled up into that indicator and the time stamp of the newest alert.

In the table view, click the Alert tab to view all alerts associated with entities in the current view of the topology. The Alert tab contains alerts as a hierarchical list. For configuration analysis alerts, the top level provides information about which policy was violated during a particular analyzer run. The next level indicates the number of times that the policy was violated. This list contains one node for each policy violation. The Alert tab only lists the policy violations that contain at least one of the affected entities that are displayed in the graphical view. The remaining policy violations are filtered. The only exception is the Overview panel where the Alert tab lists all of the policy violations. Click each violation to view a list of affected entities. The parent entities for these affected entities are displayed in a separate column. You can resort the alert list by clicking the headers in the table. For instance, by clicking the header above the entity the alert table is reorganized into a hierarchical list sorted by entity, then by alert type. If no alert was generated for an entity, the Alert column is blank.

To find an entity associated with an alert in the graphical view, click the alert below the Alert tab. The corresponding entity and any associated entities are highlighted (but not selected) in the graphical view. The highlighting is purple.

For more information, see "Viewing alerts" on page 398.

## Creating custom groups

This topic describes how custom groups can be created based on criteria that is useful and relevant to you.

Custom groups are visible in both the graphical view and the table view. You can create custom groups using any of the following criteria. The grouping criterion is displayed in parentheses in the title bar of the group box.

For example, using the topology viewer, you can organize a collection of computers into groups based on the type of operating system (for example, Windows, Linux and so on). Using a UDP, you could also create subgroups for each operating system-based group of computers (for example, the Windows group could contain subgroups for Windows 2000 and Windows NT). Organization into groups can also be based on physical location of associated devices; for example, devices in Lab A could be organized in one group, devices in Lab B could be organized in another group, and so on. To create these types of groups, use UDPs (for example, a text string of "Lab A"). To set UDPs use the pop-up menu item, Launch Detail Panel.

If you specify grouping by UDP, the groups are sorted in alphanumeric order, from left to right, in the graphical view. If the grouping is by health status, groups are

sorted, from left to right, in the following order: missing, critical, warning, normal and unknown. You can also specify the arrangement of entities within a group box based on health status and label.

To group entities, follow this procedure:

1. Right-click a group and select **Group By**.
2. Select a group-by option. Options differ depending upon the class of the group.

#### **Single Group**

All entities are included in one group.

#### **Health**

Groups are based on health status (for example, normal or warning).

**Model** The model of the storage subsystems server (for example, DS8000). This option is available for subsystem groups only.

#### **OS Type**

Groups are based on operating systems. This option is available for computer groups only.

#### **RAID level**

Groups are based on RAID level. This option is available only for pool groups at the **L2:Subsystems** level.

**Fabric** This option is available for switch groups only.

#### **User-defined properties (UDPs)**

UDPs are characteristics that you define for an entity. For example, you could create a UDP for several entities using the text string "new installation." All the entities with that UDP would be grouped together. The topology viewer provides three UDPs so that you can group resources in three different ways, depending on your needs. UDPs are not available for fabric groups.

## **Defining properties for grouping entities**

This topic describes how User-defined properties (UDPs) are useful for describing specific properties of an entity or a set of entities.

You can also use UDPs to group entities with similar characteristics. The topology viewer lets you define three UDPs for an entity. UDPs are not available for fabric groups. To define one or more UDPs for an entity, follow this procedure:

1. Right-click an entity and select **Launch Detail Panel**.
2. Type a property in the **User Defined Property 1** field (for example, a specific operating system, location of computers, and so on).
3. Optional: Type additional properties in the **User Defined Property 2** and **User Defined Property 3** fields, and then click **Set**.
4. Type a description of the element manager and click **OK**.
5. Click **Launch**.

## **Displaying data paths from one endpoint in the topology to another**

This topic describes how to view connectivity for selected entities in the topology using Data Path Explorer.

## Data Path Explorer Overview

The Data Path Explorer allows you to follow connectivity between a host disk or computer and a storage subsystem. The Data Path Explorer also provides SAN Volume Controller support. Using Data Path Explorer you can follow the connectivity from a computer to SAN Volume Controller, and from an SAN Volume Controller to a storage subsystem.

Viewing a single data path allows you to monitor performance status and pinpoint weaknesses without navigating the many entities typically seen in Topology Viewer.

Data Path Explorer shows only those entities that are part of the selected data path, and only certain types of entities, such as host, switch, host bus adapter (HBA), port, storage subsystem, and storage volume.

Data Path Explorer functions like the rest of the Topology Viewer, with key variations.

**Note:** The shortcut to Data Path Explorer is not available in the Configuration History topology view.

## Before Using Data Path Explorer

Different amounts of information can be shown by the Data Path Explorer, depending on which agents are deployed throughout the fabric.

To use Data Path Explorer for debugging performance problems, follow these steps to ensure that the information displayed in the Data Path Explorer view is current:

1. Enable Performance overlays. To see the overlays, go to topology **Global Settings > Active Overlays**: and select the **Performance** checkbox.
2. Ensure I/O is running on the host machine.
3. Create a threshold alert (see “Creating a performance threshold within an alert” on page 150).
4. Create and run performance monitors on the switches and storage subsystems of interest. See “Monitoring switch performance” on page 375 or “Monitoring storage-subsystem performance” on page 376 for detailed descriptions on creating and running performance monitors.
5. Run a discovery job for the CIM agents and Storage Resource agents or Data agents.
6. Run a probe job for the subsystem, fabric, and Storage Resource agents or Data agents.
7. Run a performance monitoring job for the subsystem and switches if you want to see performance information.

See the section about planning for the Data Path Explorer view in the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* for more information on configuring data path explorer.

## Graphical View

When you open a data path view, three tiles display in the graphical view: Initiator Entities, Fabrics, and Target Entities. These contain all the entities in all the data



paths related to the selected host or storage entity.

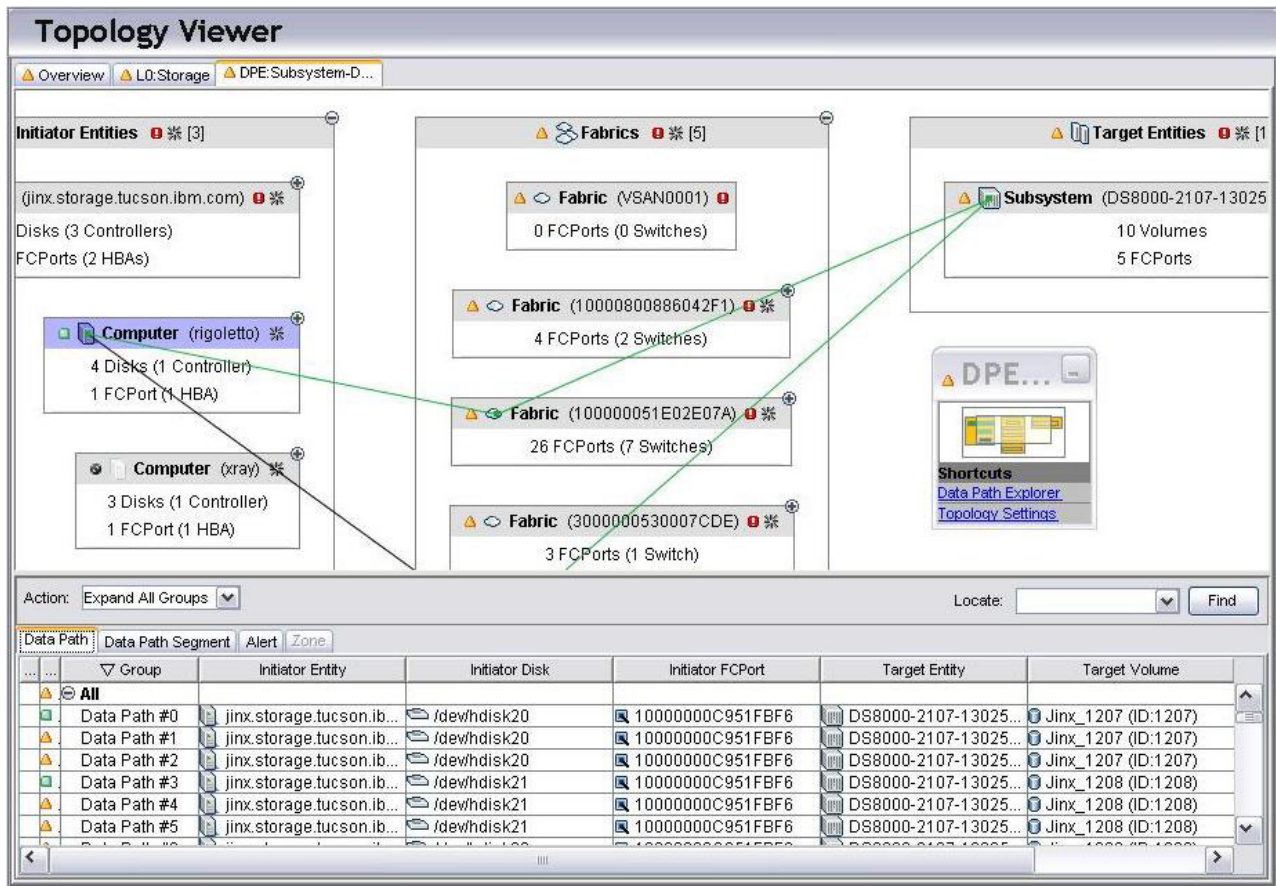


Figure 40. Data Path Explorer view

- Selecting an entity in any of the tile groups at any level displays connector lines showing the data paths the entity is involved in.
- Each tile behaves like an active topology viewer tab, allows grouping and regrouping, expanding and collapsing of groups and so forth. Each tile can be expanded to view the lower level details such as the disks, ports and volumes involved in a data path.

### Initiator Entities

The Initiator Entities tile contains the computers and related volumes, disks, and HBA ports that have valid data paths to the storage system. Expanded computer tiles display each computer, its health status icons, and a count of the entities within the computer that are part of the valid data paths. Expanded computer tiles show the controllers and disks, HBA ports, and volumes on the computer that are part of valid data paths. The initiator entities tile can also contain SAN Volume Controllers. An SAN Volume Controller is a storage virtualization device and can act as an initiator entity. Expanded SAN Volume Controller tiles show the SAN Volume Controller managed disks and ports that are part of valid data paths.

### Fabrics

The Fabrics tile contains the fabrics entities, mainly switches and ports, through which valid data paths for the selected endpoint traverse. A count of the switches and switch ports with health status icons and labels are

shown in the group title. Expanded Fabric tiles display a view of the switches to the port level. Only switches and ports with valid data paths traversing through them are displayed.

### **Target Entities**

The Target Entities tile contains the endpoint storage systems in data paths for the selected endpoint. Expanded Target tiles display each storage subsystem, its status icons, a count of entities, and the number of subsystem front-end ports and volumes involved. When these are expanded, the associated port and volume tiles display. These can also be expanded to show individual entities.

### **Table View**

The data path table displays below the graphical view. Each row in the table view shows a valid data path for the selected endpoint.

Display a single data path in the graphical view by selecting a row in the data path table.

The data path table has two tabs: Data Path and Data Path Segment.

#### **Data Path tab**

This is the default view. This tab shows each data path associated with the entity for which the data path was opened.

- Select an entity in the graphical view to display individual data paths in the table.
- The health status shown for each data path is aggregated.
- The data path table can be sorted by any of the column headers.

See “Columns in Data Path tab” on page 542.

#### **Data Path Segment tab**

This tab displays the segments in a data path. One segment is listed per row.

There are two ways to populate the Data Path Segment table:

- Select a data path in the Data Path tab. The Segment tab then lists the segments in the selected data path.
- Click an entity in the graphical view. The Segment tab lists the segments in all of the data paths the entity is involved in. These are identified by data path number (0, 1, 4, and so on).

See “Columns in Data Path Segment tab” on page 543.

**Related reference:**

“Opening a Data Path View”

Data Path Explorer provides an end-to-end topology view of the storage environment from a host disk to a storage volume. It displays the different data paths (SAN access paths or I/O paths) for each host disk.

“Displaying an entity's data paths” on page 542

You can view the data paths associated with any entity by selecting the entity in the graphical view. Lines connect the related entities in the data path or paths.

“Example use cases for Data Path Explorer” on page 544

The following examples show some of the use scenarios for Data Path Explorer:

**Opening a Data Path View**

Data Path Explorer provides an end-to-end topology view of the storage environment from a host disk to a storage volume. It displays the different data paths (SAN access paths or I/O paths) for each host disk.

You can open a data path view on any of the following endpoints:

**host** The host (computer) must have at least one disk that is assigned from a subsystem.

**subsystem**

The subsystem must have at least one volume assigned to a host.

**volume**

The volume must be assigned to a computer, and the fabric must be zoned so that the computer can see the volume.

**disk** The disk must be from a subsystem. Note that the disk can be host disk or SAN Volume Controller managed disk.

Open a data path explorer view in one of the following ways:

- Right-click an endpoint, and then select Open Data path view.
- Select an endpoint, and then select Open Data path view from the Action menu.
- Select an endpoint, and then select the Data Path Explorer shortcut in the minimap
- For multiple entities: Select one or more endpoints of the same type (for example, one or more storage volumes), and then select Open Data Path View. The data paths display for all of the endpoints selected. If you select an SAN Volume Controller entity, all other entities you select must be SAN Volume Controller entities.

**Note:**

1. Opening the Data Path Explorer view for a subsystem (including an SAN Volume Controller) that has several associated data paths can consume extensive system resources. Select volumes or disks instead, and then open the Data Path Explorer view.
2. If the endpoint is an SAN Volume Controller, two data path views are opened because the SAN Volume Controller can act as both a initiator and target entity. The first view displays the data paths associated with the SAN Volume Controller acting as a storage subsystem or a target entity. The second view displays the data paths associated with the SAN Volume Controller acting as an initiator entity.

When a data path view is opened, three tiles display: the Initiator Entities, Fabric, and Target Entities tiles. These tiles contain all the associated entities in the data

paths. Select an entity within one of the tiles to see the data paths associated with it. You may need to further expand the tiles.

If you select an entity or a group of entities that do not have a valid data path, a dialog box displays with the following message: "No available DataPath for the selected entities."

The Data Path Table lists each data path (one per row).

To return to a normal topology view, right-click an entity in the graphical or table view and select **Open Detail View** from the pop-up menu.

### **Displaying an entity's data paths**

You can view the data paths associated with any entity by selecting the entity in the graphical view. Lines connect the related entities in the data path or paths.

- In the table view-Data Path tab, the selected entity's data paths are highlighted.
- The table view-Data Path Segment tab displays the data paths belonging to the selected entity and the individual segments that make up the data path.

### **Multiple paths**

IBM Tivoli Storage Productivity Center shows the multiple paths from the disks to the storage subsystem volumes in the Data Path Explorer view. IBM Subsystem Device Driver (SDD) software allows the operating system to use multiple paths (that is, different initiator and target port pairs) through the SAN for storage access. On the L2 Computer View and the Data Path Explorer view, the disks with multiple paths managed by SDD are grouped in a controller representing the SDD driver.

For a Windows computer with multiple paths managed by SDD to subsystem volumes, a controller representing the SDD driver is displayed, and this controller contains the disks managed by the SDD.

For a non-Windows computer with multiple paths managed by SDD to subsystem volumes, Data Path Explorer displays a controller representing the SDD driver, and this controller contains the vpath disks. A separate disk (for example, hdisk on AIX) is created by the operating system on the non-Windows computer for each path to the storage volume. These disks are grouped by the controller that contains the initiator port present in the path.

The Linux SDD driver is no longer available. Starting with Red Hat Enterprise Linux Version 5 and SUSE Linux Enterprise Server Version 10, only the DM\_Multipath is available. The DM\_Multipath is supported by SAN Planner. The SDD driver is no longer supported on HP, starting with HP-UX 11i v3 with Itanium®. HP-UX 11i v3 provides a native multipathing feature which is supported by SAN Planner.

#### **Related reference:**

"Columns in Data Path tab"

When an entity in the Data Path Explorer graphical view is selected, its data paths display in the Data Path tab.

"Columns in Data Path Segment tab" on page 543

The Data Path Explorer Segment tab shows each segment in a data path.

### **Columns in Data Path tab**

When an entity in the Data Path Explorer graphical view is selected, its data paths display in the Data Path tab.

The Data Path tab columns display the following information:

**Located**

This column is associated with the Locate function in the Topology Viewer. If search text is present in a row, the Locate column displays a magnifier icon.

**Health (status)**

The aggregated status of the data path.

**Group** The number of the data path.

**Initiator Entity**

The name of the computer or SAN Volume Controller in the data path.

**Initiator Disk**

The computer disk or SAN Volume Controller managed disk in the data path.

**Initiator FCPort**

The computer Fibre Channel port in the data path. Data in this column comes from the HBA port entity in the Computer tile. If an SAN Volume Controller is the initiator entity, this column contains the SAN Volume Controller fibre channel port in the data path.

**Target Entity**

The storage subsystem in the data path.

**Target Volume**

The storage volume in the data path. An asterisk indicates that this is a space-efficient volume or vdisk.

**Target FCPort**

The storage Fibre Channel port in the data path.

**Columns in Data Path Segment tab**

The Data Path Explorer Segment tab shows each segment in a data path.

Each segment has two endpoints: the beginning device and the ending device connected by a communication link. Each endpoint resides in a parent entity, which can be a computer, fabric, or subsystem.

The Data Path Segment tab columns display the following information:

**Located**

This column is associated with the Locate function in the Topology Viewer. If search text is present in a row, the Locate column displays a magnifier icon.

**Segment Health**

The health status of the segment is an aggregation of From Endpoint Health, To Endpoint Health, From Entity Health, To Entity Health, and Connection Health.

**Label** The number of the data path and its segments. The data path number corresponds to the numbering on the Data Path tab. Segment numbering advances from left to right, from computer tile through fabric tile to subsystem tile.

**From Entity Health**

The health status of the parent entity beginning the segment.

**From Entity Performance**

The performance status of the parent entity that begins the segment. This column is displayed only when the performance overlay is enabled.

**From Entity**

The name of the computer, fabric, or subsystem that is the parent of the beginning device (endpoint) in the segment.

**From Endpoint Health**

The health status of the endpoint device that begins the segment.

**From Endpoint Performance**

The performance status of the endpoint device that begins the segment. This column is displayed only when the performance overlay is enabled.

**From Endpoint**

The name of the endpoint device that begins the segment.

**To Entity Health**

The health status of the parent entity that ends a segment.

**To Entity Performance**

The performance status of the parent entity that ends the segment. This column is displayed only when the performance overlay is enabled.

**To Entity**

The name of the computer, fabric, or subsystem that is the parent of the ending device (endpoint) in the segment.

**To Endpoint Health**

The health status of the endpoint device that ends the segment.

**To Endpoint Performance**

The performance status of the endpoint device that ends the segment. This column is displayed only when the performance overlay is enabled.

**To Endpoint**

The name of the endpoint device that ends the segment.

**Connection Health**

The health status of the connection between the endpoint device that begins the segment (From Endpoint) and the endpoint device that ends the segment (To Endpoint). The Connection Health is displayed as N/A if the health status data for the connection is not available.

**Example use cases for Data Path Explorer**

The following examples show some of the use scenarios for Data Path Explorer:

- An application on a host is running slowly, which causes you to question the health of the I/O path for that application. Are there component level performance problems that might cause the slow application response?
- Disk utilization on the storage subsystem is high, triggering performance alerts. You decide to see which hosts or applications, residing on separate host disks, are causing the problem.
- You suspect that the I/O paths for two applications (on separate host disks) are in conflict with each other. They could be sharing a common component (such as a switch). After viewing the I/O paths for these two applications, you make the required zoning or connectivity change to alleviate the problem.
- I/O path information can be used do load balancing and storage planning. For example, a host adapter port on a host has I/O from three host disks passing

through it, but a second port has I/O from ten host disks (and has more I/O). This information can be used to provision a new host disk using the first host adapter port.

**Note:** Application to host disk mapping is not shown in the Topology Viewer. This mapping must be done outside IBM Tivoli Storage Productivity Center. However, Tivoli Storage Productivity Center can provide mapping from database applications to host disks if configured to do so.

### Planning for Data Path Explorer view

This topic provides information on planning for the Data Path Explorer view in the topology viewer.

The Data Path Explorer is a new type of view in the topology viewer. Data path explorer combines several of the usual topology views into a visualization that allows users to follow connectivity from one end point in the topology view to another. Storage administrators can use this view to debug connectivity and performance problems in the storage environment. Data Path Explorer displays the different data paths (SAN access paths or I/O paths) for each host disk. This allows you to use the Data Path Explorer to check the health and performance statistics of the storage components along the I/O path.

To use Data Path Explorer for debugging performance problems, follow these steps to ensure that the information displayed in the Data Path Explorer view is current:

1. Enable Performance overlays. To see the overlays, go to topology **Global Settings > Active Overlays**: and select the **Performance** checkbox.
2. Ensure I/O is running on the host machine.
3. Create a threshold alert (see “Creating a performance threshold within an alert” on page 150).
4. Create and run performance monitors on the switches and storage subsystems of interest. See “Monitoring switch performance” on page 375 or “Monitoring storage-subsystem performance” on page 376 for detailed descriptions on creating and running performance monitors.
5. Run a discovery job for the CIM agents and Storage Resource agents or Data agents.
6. Run a probe job for the subsystem, fabric, and Storage Resource agents or Data agents.
7. Run a performance monitoring job for the subsystem and switches if you want to see performance information.

Different agents play different roles when discovering information and affect what can be shown in the Topology Viewer. The Data Path Explorer view is also subject to those limitations. Different amounts of information can be shown by the Data Path Explorer view depending on the agents deployed throughout the fabric.

To display data paths from a host to a subsystem in this view, it is necessary to have the following agents installed and configured:

- CIM agent to monitor the storage subsystem
- Data agent to monitor the host
- In-band Fabric agent to monitor the fabric. Also a CIM agent for the fabric to collect fabric performance information.

To show the host disk assignment from a given storage subsystem, IBM Tivoli Storage Productivity Center Data Server needs access to the subsystem's CIM

agents (SMI-S agent). In the case of the subsystem, the data path for the subsystem is defined by the host masking or mapping information maintained by the CIM agent.

Analysis on data paths involving SAN Volume Controllers requires one or more CIM agents pointing to the SAN Volume Controller and the back-end system.

To show accurate information, it is important to have the most current information from the CIM agents, Fabric agents, and Data agents. You must probe all the components within the data path (or a set of agents strategically placed throughout the fabric) before using the Data Path Explorer.

To show physical fabric connectivity, you must have an in-band Fabric agent, out-of-band Fabric agent, or a CIM agent for the fabric.

To show performance overlay in the Data Path Explorer view, you must first create and run performance monitors on the switches and storage subsystems of interest. To see the overlays you must also go to **Topology > Global Settings > ActiveOverlays** and select the **Performance** check box. See “Collecting performance data” on page 375 for detailed descriptions on creating and running performance monitors.

To show zone overlay, you need an in-band Fabric agent connected to the fabric along the data paths defined or a CIM agent for the fabric. To see the overlays you must also go to topology Global Settings and select Show zone tab check box. **Topology > Global Settings** and select the **Show zone tab** check box.

If you want to launch the Data Path Explorer view for a host, that host must have a Data agent running on it.

These are the general setup requirements for using the Data Path Explorer:

1. Install and configure the SMI-S agents (CIM agents) for the subsystems (including SAN Volume Controllers) and fabric to collect subsystem and fabric data. Contact your vendor for information on installing and configuring the CIM agent for the subsystem or fabric.
2. Install and configure the in-band Fabric agent, or out-of-band Fabric agent to monitor a given fabric.
3. Install the Data agent on the hosts.

## Displaying membership in zones and zone sets

This topic describes how the topology viewer allows you to examine zones and zone sets.

Zone information is shown in the **Zone** tab in the table view. The **Zone** tab is disabled by default. Selecting a zone or zone set in the **Zone** tab highlights zone and zone set members in the graphical view. Ports and entities that are members of the selected zone or zone set are indicated in the graphical view by highlighting. You can select highlighted entities to show details (for example, connections).

The entities listed in the **Zone** tab are only those entities that are in the current view. However, zone or zone-set selections apply to all currently open views in the topology viewer. Therefore, switching to another view might indicate other entities in the selected zone.

**Note:** You cannot use the **Zone** tab to change membership status of entities.



To activate the **Zone** tab, follow this procedure:

1. Right-click in the graphical view and select **Global Settings**.
2. In the Settings window, select **Show zone tab** and click **OK**.

## Displaying detailed information about entities

This topic describes how to view detailed information about entities using the topology viewer.

There are two ways to display information about entities:

1. Right-click the entity and select **Launch Detail Panel** to display detailed information about an entity (for example, the operating system of a computer, the used and unused capacity of a storage subsystem)
2. Right-click the entity or entities and select **Reports**. From the dialog displayed, you can choose which report to create. The reports you can choose depend on which entities you selected. Asset reports are available only if you selected a single entity. Select a report and click **Create Report**

## Displaying storage resource connections

This topic describes how physical and logical connections among entities in the environment can be displayed using the topology viewer.

Connections are shown in both the graphical and table views. In the graphical view, connections are represented by lines between entities or groups. In the table view, a **Connections** tab lists both ends of a connection.

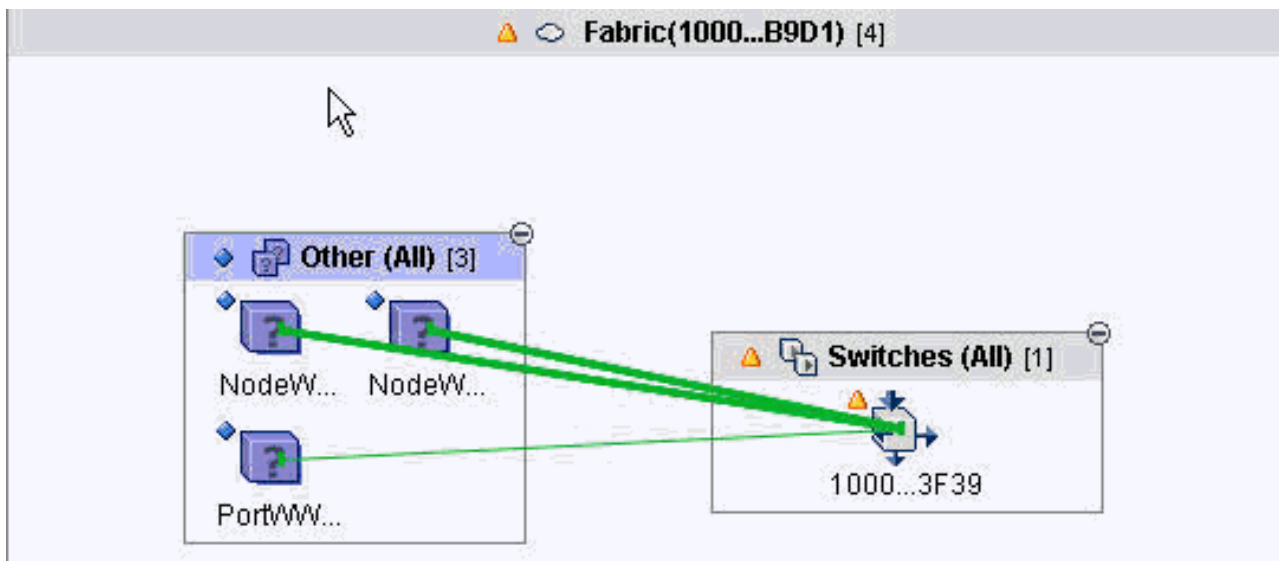


Figure 41. Connections

In the graphical view, the color of a connecting line between entities indicates the status of a connection. The color of the line between groups, individual entities, or between an individual entity and a group can be green (normal), yellow (warning), red (critical), black (missing), or grey (status unavailable).

**Note:** When you have a host with two HBAs or FC ports that are both in the same switch and you switch the ports between them, the connection will be shown with

two yellow (warning) lines. This is because you have not reset the history retention setting. Until the settings are reset, the connecting lines will remain yellow.

Connections between two individual entities are represented by a thin line. Connections between groups or between a group and an individual entity are represented by a thick line. A connection to a collapsed group can represent either a single connection or several connections.

To show connections, follow this procedure:

1. Right-click a blank space in background of the graphical view and select **Global Settings**.
2. In the Settings window, select **Show Connections** and click **OK**.

## Filtering entities based on health status

This topic describes how to quickly narrow the health status of normal or non-normal entities in your environment.

In views with many entities and connections, the display can become very cluttered and can make it difficult to determine where problems are occurring. You can filter entities to display only those entities whose health status is non-normal (critical, warning, missing, or unknown).

When you apply filtering, you can also choose to display entities with a health status of normal but that are connected to entities whose health status is either critical or warning. To filter entities by health status:

1. Right-click in the background area of the graphical view and select **Global Settings**.
2. In the Settings window, select **Hide Normal Entities** and, optionally, **Unless connected to non-normal**, and click **OK**.

## Labeling entities

This topic describes how you can specify or change label entities.

In both the graphical and table views, entities are identified by labels (for example, the name of a computer). You can specify or change these labels. To specify or change the existing label for an entity, follow this procedure:

1. Right-click the entity and select **Launch Detail Panel**.
2. Type a label for the entity in the **Label** field.
3. Click the save icon in the menu bar to save the label, and then close the window.

## Starting element managers

This topic describes how to start element managers for the entities that are defined in IBM Tivoli Storage Productivity Center.

If element managers on a system, tape library, switch, or other entity have been defined, you can start the element managers to access vendor-provided management tools.

To start an element manager for an entity, right-click the entity. A menu that contains launch options is displayed. The options that are displayed depend on the element managers that you have configured on the Element Manager page of the Configuration Utility. The following options are displayed on this menu:

- **Launch in context defined application.** This menu item is displayed for each application that is defined by a launch descriptor file. Tivoli Storage Productivity Center provides launch descriptor files for the following applications:
  - System Storage DS8000 Storage Manager
  - The management console for the storage systems that are listed under **Management GUI**.
  - Tivoli Storage Productivity Center for Replication

If the corresponding launch descriptor files are found for the applications, the following menu items are displayed:

#### **DS8000 Storage Manager**

Click this menu item to start DS8000 Storage Manager.

#### **Management GUI**

Click this menu item to start the management console for the following storage systems:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

#### **Replication Storage Systems Overview**

Click this menu item to start Tivoli Storage Productivity Center for Replication. The window that is opened in the Tivoli Storage Productivity Center for Replication GUI depends on the context for the entity that you selected. For example, if you right-click a storage device on the **Element Manager** tab of the Configuration Utility and click **Replication Storage Systems Overview**, the Tivoli Storage Productivity Center for Replication Storage Systems window opens.

If the launch descriptor file cannot be found, the **Launch file items error** menu item is displayed. Click this item to view a description of the problem. The **Launch file items error** menu item is also displayed in situations such as the IBM Tivoli Integrated Portal server is not installed or is not accessible, the Tivoli Storage Productivity Center Data server is not accessible, or another system error occurred.

- **External Tools.** This menu item shows any external tools that are defined in the Configuration Utility. Examples of external tools include Web-based user interfaces and local executable files. For example, you might add the address of a website that you frequently access or an application that you often use such as a text editor or a spreadsheet.
- **Element Manager.** This menu item shows any element managers discovered for the device by a Storage Management Initiative Specification (SMI-S) query. To avoid menu item duplication, this item is not displayed when the associated element manager is defined by a launch descriptor file. In this situation, the element manager is displayed as a separate menu item.
- **User Defined Tools.** This menu item shows any user-defined element managers that are defined in the Configuration Utility. These element managers differ from the element managers in the Element Manager menu because they are user-defined and are not discovered by a SMI-S query.

If the selected entity has no element manager, these options are not available in the menu.

## Launching SAN Planner

This topic describes how to launch the SAN Planner wizard for an entity from the topology viewer. SAN Planner assists end-to-end planning for SAN components.

Right-click a computer, fabric, or storage subsystem and select **SAN Planner > New Plan**. A new SAN Planner wizard introduction page is displayed. When you have made your selections on the Select Planning Task and Select Provision Plan Details pages, the element you selected is added to input resources on the Specify Plan Content page. To provision replication for an existing session associated with a storage resource group, right-click the storage resource group, select **SAN Planner**, and then select the session you want to work with.

### Related concepts:

“SAN Planner overview” on page 417

SAN Planner assists the user in end-to-end planning involving fabrics, hosts, storage controllers, storage pools, volumes, paths, ports, zones, zone sets, storage resource groups (SRGs), and replication. Once a plan is made, the user can select to have the plan implemented by SAN Planner.

## Starting Storage Optimizer

This topic describes how to start the Storage Optimizer from the topology viewer. The Storage Optimizer uses data in the Tivoli Storage Productivity Center database to analyze your storage subsystems to identify performance bottlenecks, and recommend changes to improve performance.

Right-click one or more storage subsystems and select **Launch Optimizer**. A new Storage Optimizer analysis definition is created, and the storage subsystems you selected are added to the **Select Storage** pane.

### Related concepts:

“Storage Optimizer overview” on page 435

The Storage Optimizer uses data in the IBM Tivoli Storage Productivity Center database to analyze your storage subsystems to identify performance bottlenecks, and recommend changes to improve performance. This topic lists the supported subsystems and describes the general steps for using Storage Optimizer.

## Locating storage resources in the topology

This topic describes how you can search for a particular entity using the entity name or a keyword (for example, a host name or operating system).

The open view (for example, L0:Other) will be searched for the string you enter. The results of a locate action are shown by highlighting.



Figure 42. Highlighted entity

To locate an entity, type a search string in the **Locate** field in the right corner of the table view and click **Find**. The current view (for example, L0:Other) will be searched for the string you enter. The locate function does not search the entire topology, only the currently open view.

When the entity is found, it will be displayed in both the graphical and table views. If the located entity is inside a collapsed group in the topology, you will not be able to see that entity until you expand that group. However, the entity will be highlighted in the table view. If the entity is in a collapsed group in the table view, that group will automatically be expanded and the highlighted table entry will receive focus. You can then click that table entry to see the located entity in the topology view.

If you specify a Locate criterion that applies to more than one entity in a view, all the applicable entities are displayed.

## Managing Storage Subsystems

This topic describes how to display the Storage Subsystems page, where you can perform storage-subsystem management tasks.

Right-click a storage subsystem and select **Manage Storage Subsystems**.

## Managing Zone Control

This topic describes how to display the Fabric page, where you can perform zone-control tasks.

Right-click a fabric and select **Manage Zone Control**.

## Pinning entities

This topic describes how pinning entities is an easy way to identify selected entities for future reference or quick access.

Pinned entities display a small flag next to the entity. This flag is propagated to all views that contain that entity. This is useful for marking entities for various reasons, such as a reminder to look at this entity later and for upward propagation. For example, the user may pin a computer in an L0 or L1 Computers view and then quickly wish to see which fabric it belongs to by opening an L0 Fabrics view and looking at which fabric group contains the pinned entity - indicated by the small flag. Pinned objects remain visible as you change views in the topology viewer and when you log off and log back in. Pin lists are stored on a per user basis. This means that different users can have separate pin lists. Pinned objects are indicated by a pin icon in the **Pinned** column of a table or by a flag icon next to the entity in the graphical view.



Figure 43. Pinned entity

To pin an entity, right-click the entity in the graphical view and select **Pin**. To unpin an entity, right-click the entity in the graphical view and select **Unpin**. The same function is available in the **Action** list in the table view.

Pinned entities are visible in groups at all levels and in all views regardless whether the group is expanded or collapsed. You can only pin individual entities; you cannot pin groups (for example, a group of computers).

## Printing information from the topology viewer

You can print information displayed in the topology viewer on a single page for reference.

Printing capability is available for the graphical view but is not available in the tabular view. To print the information displayed in the topology viewer, right-click the graphical view and select **Print** or **Print Preview**. The content is scaled to fit onto one page if necessary.

## Refreshing views

This topic describes how you can refresh the topology viewer to view the most up-to-date details about your environment.

To provide you with the most up-to-date information, the topology viewer automatically refreshes all open views, including graphics and tables, every five minutes from the IBM Tivoli Storage Productivity Center database. You can also refresh views manually or change the default refresh rate. Opening a new view automatically shows the most current information. To manually refresh views or change the refresh rate, right-click in the graphical view or click the arrow in the **Action** field in the table view and choose one of the following options:

- **Refresh View:** Updates the view that is currently displayed.
- **Refresh All Views:** Updates all views that are currently open. You can also use this option to remove the alert indicators from the devices in the topology view.
- **Refresh Settings:** Opens the **Refresh Rate Settings** window, in which you can select a new automatic refresh rate (specified in minutes).

## Removing entities from the database

Computers in the missing state were once in the topology viewer, but were not discovered by IBM Tivoli Storage Productivity Center during the latest refresh cycle. You might want to remove them from the view and the database in order to make it easier to work with other resources.

To remove a computer from the topology viewer (both graphical and tabular views) and the database, you first need to uninstall the agents on that system.

**Note:** This action is available only for entities that are in the missing state.

1. See *Uninstalling the agents* in the *IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* for instructions on uninstalling the agents.
2. In the topology viewer *graphical* view, right-click the computer entity and select **Remove From Database..**
3. In the topology viewer *tabular* view, right-click the computer entity and select **Remove From Database.**

## Removing deleted entities

When a discovered entity (fabric, subsystem, switch, or volume) is deleted using another application, such as an element manager or other management application, the next discovery or probe of the topology viewer displays a missing entity. The status of missing entities is aggregated and propagated up through the device hierarchy.

It can be difficult to determine if there are other lower-level entities that have "missing" status, but have not been intentionally removed by a management application. The Removed Resource Retention setting in the topology viewer can be set to a low value, such as 0 or 1 day to remove the missing entity from the display. This removes the missing entities from the display the next time the Removed Resource Retention process runs.

You can also remove deleted entities from the topology viewer using the following procedure:

1. Highlight the intentionally deleted entities in the topology view and select **Remove Resource** in the **Action** list in the tabular view.

**Note:** Not all entity classes can be removed from the topology using the **Action** list. For those entities, use the Removed Resource Retention setting.

2. A prompt asks you if you are sure you want to remove the resource from the database. Click **Yes**.
3. To refresh the view, click **Refresh View** in the **Action** list. The missing entities will no longer be displayed in the topology viewer.

## Selecting entities in graphical and table views

This topic describes how to select entities in the graphical and table views of the topology viewer.

To select an entity in the graphical view, click the entity image. In the table view, click the name of the entity. To make multiple selections, press CTRL and click.

## Showing Alerts

This topic describes how to display alert logs for entities in the topology viewer.

You can display alert logs for:

- computers
- fabrics
- storage subsystems
- tape libraries
- switches

All alerts for the entity-type you select will be displayed, even if no alerts exist for the entity.

Right-click an entity and select **Show <entity type> Alerts**.

For example:

Right-click a computer and select **Show Computer Alerts**.

## Showing the monitoring status of entities

View how Tivoli Storage Productivity Center is monitoring a storage entity that is displayed in the topology viewer.

Display all the data collection schedules, runs, and jobs that are related to a storage entity in the topology viewer.

To view monitoring status, right-click an entity and select **Show Monitoring Status**. For example, right-click a computer and select **Show Monitoring Status**.

When you select this action, the Job Management panel is displayed. Use this panel to view information about how a storage entity is monitored by Tivoli Storage Productivity Center. For storage resource groups, when you click this option the schedules for all the storage entities in that group are displayed.

### Related reference:

“Managing schedules, jobs, and runs” on page 295

Use the Job Management panel to view and manage all the schedules, runs, and jobs related to the storage devices that are monitored by IBM Tivoli Storage Productivity Center.

## Working with tape libraries in the topology viewer

Perform storage management functions on your tape library using the topology viewer.

The topology viewer displays information about the discovered tape libraries and tape library groups. The topology viewer has two views:

- Topology view
- Tabular view

Each view shows information about the discovered storage devices in your enterprise. Changes to one view are automatically displayed in the other view, so that both views are always synchronized.

To see tape library information in the topology viewer, click the **Storage** node in the **IBM Tivoli Storage Productivity Center > Topology** node in the console. The level zero (L0) view of the storage topology view is displayed. The view includes the tape libraries and tape library groups that have been discovered.

### Viewing tape libraries in the topology view

This topic describes how you can see information about your tape libraries and tape library groups in the topology view.

Clicking a tape library icon in the topology view highlights (selects) that tape library in both the topology view and the tabular view. Right-click a tape library to display a pop-up menu that lists actions that you can perform on the library, as well as other general actions that you can perform. The actions specific to tape libraries are:

#### Open Detail View

Opens the next level of detail for the selected topology object. For tape libraries, the view hierarchy is

- L0:Storage- overall view of the discovered storage subsystems and tape libraries
- L1:Tape Libraries- tape library groups



- L2:Tape Library- view of an individual tape library and its related devices and connections

#### **Launch Detail Panel**

Displays the detail panel for the tape library. The detail panel shows detailed information about the tape library, such as the name of the tape library vendor, the serial number, status, owner/administrator contact information, and properties that can be defined by the user. It also has **Launch** button that will launch the element manager for the tape library.

#### **Launch Element Manager**

Launches the element manager application for the selected tape library. The element manager is an application or Web site provided by the tape library vendor for performing management tasks for the selected tape library. For example, the element manager for an IBM 3584 tape library is the IBM Tivoli Storage UltraScalable Tape Library Specialist.

**Pin** Pins the selected tape library. When a topology object is pinned, it remains visible even after its respective group or view level has been collapsed. A small blue flag icon to the right of the object's icon indicates that the object is pinned. Once an object has been pinned, this action will change to **Unpin** in the pop-up menu for the object.

As you add more storage devices to your enterprise, your topology view can become very crowded and increasingly difficult to navigate. You can use the **minimap** in the upper right corner of the topology view to quickly move around the view and locate the objects you need to work with.

### **Viewing tape libraries in the tabular view**

This topic describes how you can see information about your tape libraries and tape library groups in a tabular view.

The tabular view presents the same information as the topology view. Changes to the tabular view are immediately reflected in the topology view, so that both views are always synchronized.

At the top of the tabular view are two fields:

#### **Action**

Lists the actions that can be performed on the selected object in the table, as well as the more general actions, such as Expand All Groups, that can be performed. These actions are the same ones that appear in the context menu for a selected object.

**Locate** Allows you to find objects in the tabular view. This is very useful when your tabular view has many entries.

The columns in the tabular view change as you select different objects in the topology. For example, the columns available for objects in the L1:Tape Libraries view are different than the columns in the L2:Tape Library view. These columns show information about the tape libraries, such as the operational status, vendor name, and number of drives in the library.

---

## **Managing disk storage**

This section describes managing storage subsystems, virtual disks, volumes, and creating and managing alerts.

## Working with workload profiles

This topic describes creating, modifying, and deleting workload profiles.

Only storage subsystems that have been discovered and had performance data successfully collected can be used with the Volume Planner function of the SAN Planner. You can use the following products with SAN Planner:

- TotalStorage Enterprise Storage Server
- System Storage DS6000
- System Storage DS8000
- SAN Volume Controller
- Storwize V7000
- XIV system

For information about storage subsystem discovery, see “Discover storage subsystem, tape library, and fabric information with CIM agents” on page 330. For information about the Volume Planner, see “Using the SAN Planner” on page 422.

### Creating a workload profile

This topic describes how to create a workload profile. You can use the performance information from the selected storage-system volumes to generate a workload profile. In SAN Planner, you can create or reuse volumes based on the workload profile.

1. In the left pane, expand **Disk Manager > Profile Management > Workload Profiles**. In the right pane, the currently defined workload profiles are displayed.
2. Click **Create**; the Workload Profile wizard starts.
3. On the Welcome to the Workload Profile Wizard page, click **Next**.
4. On the Creation Method page, select the basis for the new workload profile. Click one of the following choices:

#### **Create profile from volume performance data**

Base the new profile on the volume-performance data that you provide.

#### **Create from an existing profile**

Base the new profile on an existing workload profile. In the **Profile Name** field, select the workload profile.

Click **Next**. If you clicked **Create from existing profile**, go to step 6. Otherwise, the Name and Description page opens.

5. On the Name and Description page, provide information about the new workload profile:
  - a. In the **Name** field, type the name of the workload profile.
  - b. Optional: In the **Description** field, type a description of the workload profile.
  - c. Click **Next**.
6. On the Select Volumes page, specify the volumes whose performance data will be used as the basis for the workload profile.
  - a. In the **Available volumes** field, select the volumes, and then move them to the **Selected volumes** field.
  - b. Click **Next**.
7. On the Peak Activity periods page, specify the time period for the performance data:

- a. To use all available performance data, click **Use all available performance data**, and go to step 8.

**Note:** Analyzing all performance data might be a lengthy process.

- b. To use performance data for a specified time period, specify the time period. In the **Start date** field, select the month, day, and year. In the **End date** field, select the month, day, and year. By default, the time period is set to the previous month. In the **Time reference** field, select one of the following options:

**Client Time**

**Server Time**

**GMT**

**Device Time**

- c. To specify that performance data is collected only during specified hours, click **Use only a portion of each day**. In the **Start time** field, select the hour and minutes. In the **End time** field, select the hour and minutes.
  - d. Click **Next**.
8. On the Review Profile Results page, review the following performance data:

#### **Data quality**

Lists a data-quality rating. This rating is based on the percentage of volumes that had valid data:

- Good (100% - 67%)
- Fair (66% - 33%)
- Poor (32% - 0%)

#### **I/O Demand**

##### **I/O rate per GB per sec**

The data-transfer rate in gigabytes per seconds (GBps).

##### **Avg. transfer size (KB)**

Default size for data transfer in kilobytes (KB). You can edit this value.

#### **Read / Write**

##### **Sequential reads (%)**

Percentage of read/write activity that was sequential reads.

##### **Sequential writes (%)**

Percentage of read/write activity that was sequential writes.

##### **Random reads (%)**

Percentage of read/write activity that was random reads.

##### **Random writes (%)**

Percentage of read/write activity that was random writes.

##### **Total (%), must be 100**

Total of the preceding percentages. This value must equal 100%. If it does not, you cannot continue with the wizard.

#### **Cache**

##### **Random read cache hits (%)**

Percentage of cache activity that was random hits to read cache.

### Random write destage (%)

Percentage of cache activity that consisted of moving data from cache to nonvolatile storage.

9. Click **Finish**.

### Deleting a workload profile

This topic describes how to delete a workload profile.

Only user-defined workload profiles can be deleted.

1. In the left pane, expand **Disk Manager > Profile Management > Workload Profiles**. A list of the currently defined workload profiles is displayed in the right pane.
2. Select the workload profile you want to delete.
3. Click **Delete**.

### Modifying a workload profile

This topic describes how to modify a workload profile.

Only user-defined workload profiles can be modified.

1. In the left pane, expand **Disk Manager > Profile Management > Workload Profiles**. A list of the currently defined workload profiles is displayed in the right pane.
2. Select a workload profile and click **Details**. The Workload Profile Properties window displays.
3. Click the **General** tab to view the name and description of the workload profile. You can modify only the description of the workload profile.
4. Click the **Performance** tab to view I/O demand, read/write, and cache performance data for the workload profile. You can modify any of the performance data for the workload profile.

### Viewing a workload profile

This topic describes how to view a workload profile.

1. In the left pane, expand **Disk Manager > Profile Management > Workload Profiles**.
2. In the right pane, select a workload profile and click **Details**. The Workload Profile Properties window displays.
3. Click the **General** tab to view the name and description of the workload profile.
4. Click the **Performance** tab to view I/O demand, read/write, and cache performance data for the workload profile.
5. Click the **Activity** tab to view peak activity settings for the workload profile. To view all data collected during a period of time, select **Use all available performance data** and specify a time reference (options are Client Time, Server Time, GMT, and Device Time) and start and end dates. To view only the data collected during a portion of each day, select **Use only a portion of each day** and specify start and end times. Click **OK** to view the peak activity.

## Working with provisioning profiles

This topic describes creating, modifying, and deleting provisioning profiles.

Only storage subsystems that have been discovered and had performance data successfully collected can be used with the storage provisioning function of the

SAN Planner. SAN Planner and storage provisioning are only available for TotalStorage Enterprise Storage Server, SAN Volume Controller, DS6000, DS8000, and IBM Storwize V7000 systems.

For information about storage subsystem discovery, see “Discover storage subsystem, tape library, and fabric information with CIM agents” on page 330. For information about SAN Planner, see “Using the SAN Planner” on page 422

## Creating a provisioning profile

This topic describes how to create a provisioning profile. Provisioning profiles are used by SAN Planner, and in general, by features that require a model for volume, multipathing, and zoning requirements.

1. In the left pane, expand **Disk Manager > Profile Management > Provisioning Profiles**. In the right pane, the currently defined provisioning profiles display.
2. Click **Create**. The Provisioning Profile wizard starts.
3. On the Welcome to the Provisioning Profile Wizard page, provide information about the new provisioning profile.
  - a. In the **Profile Name** field, type the name of the provisioning profile.
  - b. Optional: In the **Profile Description** field, type a description of the provisioning profile.
4. In the Creation Method section, choose the method that you want to use to create the profile. Select one of the following options:

### Create profile from scratch.

Click to create the profile without using an existing profile.

### Create from an existing profile.

Click to create the profile by using an existing provisioning profile as a template. Select the profile from the **Profile name** list.

Click **Next**.

5. If you want to include volume planning in the provisioning profile, use the Volume Settings page to configure and create volumes or virtual disks.
  - a. Select **Specify how the storage will be allocated and the storage performance characteristics** to enable the options on this page.
  - b. Select one of the following options:
    - **Provision Volumes**. Click to use volumes for provisioning.
    - **Provision Virtual Disks**. Click to use virtual disks for provisioning.
  - c. In the **Total Capacity** field, enter the total capacity of the volumes or virtual disks that are to be used for provisioning storage in GB.
  - d. Enter the number of volumes or virtual disks that you want to create in the **Number of Volumes/Virtual Disks** field.
  - e. From the **RAID level** list, select one of the following Redundant Array of Independent Disks (RAID) levels for the volume or virtual disk:
    - **<system selected>**. The best possible RAID level of the volume is selected based on user input and the available performance and capacity of the underlying storage systems.
    - **RAID 1**. The format type of the volume will be RAID 1 format.
    - **RAID 5**. The format type of the volume will be RAID 5 format.
    - **RAID 6**. The format type of the volume will be RAID 6 format.
    - **RAID 10**. The format type of the volume will be RAID 10 format.

- f. In the **Volume/Virtual Disk Name Prefix** field, type in a prefix to be used in the volume or virtual disk name. This value is used as the prefix for all volumes or virtual disks that are created.
- g. From the **Workload Profile** list, select a workload profile for the volume or virtual disk. You can use one of the following predefined profiles or you can create a profile from the Workload Profiles panel. A performance monitor must have been run on the storage controller to select any predefined workload profile other than **Space only**.
  - **OLTP Standard.** Use for typical online transaction processing.
  - **OLTP High.** Use for very active online transaction processing.
  - **Data Warehouse.** Use for applications with inquires into large data warehouses.
  - **Batch Sequential.** Use for batch applications involving large volumes of data.
  - **Document Archival.** Use for document archival applications.
  - **Space only.** Use to select pools based on pool space information rather than storage controller performance. Use for a storage system that has not been monitored for a performance data collection.
- h. From the **Thin Provisioning Profile** list, select a thin provisioning profile if you want to create thin volumes or virtual disks for IBM System Storage SAN Volume Controller, IBM System Storage DS8000, or IBM XIV Storage System. You can use one of the following predefined profiles or you can create a profile from the Thin Provisioning Profiles panel.
  - **Ignore thin provisioning.** A thin provisioning profile is not used.
  - **Standard.** This profile sets the over allocation limit for a pool at 200%. This percentage sets the virtual capacity at up to two times the real capacity.
  - **Conservative.** This profile sets the over allocation limit for a pool at 150%. This percentage sets the virtual capacity at up to one and a half times the real capacity.
  - **Aggressive.** This profile sets the over allocation limit for a pool at 300%. This percentage sets the virtual capacity at up to three times the real capacity.

The over allocation value identifies the limit on the virtual capacity allocated for a pool. A lower over allocation percentage is considered conservative because the virtual capacity is closer to the real capacity. A higher allocation percentage is considered aggressive because the virtual capacity is farther from the real capacity. For example, the over allocation limit that is set in the Aggressive profile is 300%. Pools that contain volumes or virtual disks that are created using this profile will be allocated up to three times the real capacity. If you use the Aggressive profile and the pool size is 15 GB, the virtual capacity allocated for the pool can be up to 45 GB.
- i. Select one of the **Solid State Disks** options to specify whether you want to use or exclude solid-state disks from provisioning:
  - **Do not use Solid State Disks.** Click to exclude solid-state disks from provisioning.
  - **Use Solid State Disks.** Click to use only solid-state disks for provisioning.

These options are disabled if you clicked **Provision virtual disks**.
- j. From the **Encryption Group** list, select an encryption group if you want new volumes to be created on extent pools that belong to that encryption group.

If there are no encryption groups available, **<none selected>** is displayed in this box. This box is disabled if you clicked **Provision virtual disks**.

- k. Select **Use unassigned volumes not involved in copy relationship** if you want to use existing unassigned volumes. Unassigned volumes that are in a copy relationship are not used. If this check box is selected, the following volumes are excluded:
  - Unassigned volumes on non-IBM storage systems.
  - Unassigned volumes on IBM storage systems that have copy relationships that were not created in IBM Tivoli Storage Productivity Center or IBM Tivoli Storage Productivity Center for Replication.
- l. Enter the number of days of performance data that you want to use for planning in the **Planning will be based on the performance data collected in the last x days** field.

Click **Next**.

6. If you want to include path planning in the provisioning profile, use this page to set multipath options. The supported multipath driver on a host is the IBM Subsystem Device Driver (SDD). The Linux SDD driver is no longer available. Starting with Red Hat Enterprise Linux Version 5 and SUSE Linux Enterprise Server Version 10, only the DM\_Multipath is available. The DM\_Multipath is supported by SAN Planner. The SDD driver is no longer supported on HP, starting with HP-UX 11i v3 with Itanium<sup>®</sup>. HP-UX 11i v3 provides a native multipathing feature which is supported by SAN Planner.
  - a. To set up the multipath options, click **Setup multipath options (if supported by the host drivers)**.
  - b. From the **Multipath Option** list, select one of the following options to specify how the driver uses the paths between the host and the storage system.
    - **Load Balancing** sends input/output on all paths.
    - **Round Robin** sends input/output on one path until a time interval that is set in an SDD setting at the host expires, then switches to another path.
    - **Fail-Over** Sends input/output on one path until a failure occurs, then switches to another path.

On HP-UX 11i v3 with Itanium<sup>®</sup>, the multipath options used by the SAN Planner are as follows:

- **Load Balancing** on HP-UX 11i v3, is implemented as Least Command Load. This option selects the path with the least number of pending I/O requests for the next I/O operation.
  - **Round Robin** on HP-UX 11i v3, distributes the I/O load equally across all active paths irrespective of the current load on each path.
  - **Fail-Over** on HP-UX 11i v3, is implemented as Preferred Path. For this policy, you specify a path to an optimized controller. This path is preferred for I/O transfers. The preferred path must be set manually.
- c. Select the **Specify number of paths** check box and enter the number of paths that you want to configure between the host and storage system. If you do not specify the number of paths in this field, the default is two paths.
  - d. Select **Use fully redundant paths (requires at least 2 fabrics)** to use the paths from host to storage systems through a minimum of two fabrics.

Click **Next**.

7. Use the Security Settings section to specify zoning between ports on the hosts and systems in a fabric.

- a. Select **Automatically change zoning to ensure hosts can see the new storage** to enable the options on this page.
  - b. Select from the **Automatically create zone** list to specify how the zones will be created. The predefined options are:
    - <auto-zone>**  
This option is the default. The plan generates using the maximum number of zones without grouping the data paths based on hosts, host ports, or host bus adapters (HBAs).
    - ...for each host**  
Creates a zone for each host.
    - ...for each HBA**  
Creates a zone for each host bus adapter.
    - ...for each host port**  
Creates a zone for each host port.
  - c. Select the **Specify maximum number of zones** check box and enter the maximum number of zones.
  - d. Select the **Maximum zone members per zone** check box and enter the maximum number of zone members per zone.
  - e. Select **No two controllers with different types should be used in the same zone** if you do not want two controllers with different types to be used in the same zone.
  - f. Select **Use active zone set** if you want any zone set that is available to be selected.
  - g. Select the **Append zone name prefix** check box and enter a prefix to be used in the zone name.
8. Click **Finish**.

## Deleting a provisioning profile

This topic describes how to delete a provisioning profile.

Only user-defined provisioning profiles can be deleted.

1. In the left pane, expand **Disk Manager > Profile Management > Provisioning Profiles**. A list of the currently defined provisioning profiles is displayed in the right pane.
2. Select the provisioning profile you want to delete.
3. Click **Delete**.

## Modifying and viewing a provisioning profile

This topic describes how to modify and view a provisioning profile.

Only user-defined provisioning profiles can be modified.

1. In the left pane, expand **Disk Manager > Profile Management > Provisioning Profiles**. A list of the currently defined provisioning profiles is displayed in the right pane.
2. Select a provisioning profile and click **Details**. The Provisioning Profile Properties window displays.
3. Click the **General** tab to view the name and description of the provisioning profile. You can modify only the description of the provisioning profile.
4. Click the **Volume Settings** tab to view or modify the volumes or virtual disks in the provisioning profile. You can modify any of the volume settings for the provisioning profile.



5. Click the **Path Settings** tab to view or modify multipath planning in the provisioning profile. You can modify any of the multipath settings for the provisioning profile.
6. Click the **Security Settings** tab to view or modify zoning between ports on the hosts and systems in a fabric. You can modify any of the security settings for the provisioning profile.

## Working with thin provisioning profiles

Learn how to create, modify, and delete thin provisioning profiles.

Only storage subsystems that have been discovered and had performance data successfully collected can be used with the volume planning function of SAN Planner. SAN Planner and the volume planning option are only available for TotalStorage Enterprise Storage Server, SAN Volume Controller, DS6000, DS8000, and IBM Storwize V7000 systems.

For information about storage subsystem discovery, see “Discover storage subsystem, tape library, and fabric information with CIM agents” on page 330. For information about the volume planning option, see “Using the SAN Planner” on page 422

### Creating a thin provisioning profile

Learn how to create a thin provisioning profile. Thin provisioning profiles serve as templates for specifying the parameters for thin provisioned volumes to be created through SAN Planner.

1. In the left pane, expand **Disk Manager > Profile Management > Thin Provisioning Profiles**. In the right pane, the currently defined thin provisioning profiles display.
2. Click **Create**. The Thin Provisioning Profile window opens.
3. On the Thin Provisioning Profile page, provide information about the new thin provisioning profile.
  - a. In the **Profile Name** field, type the name of the provisioning profile.
  - b. Optional: In the **Profile Description** field, type a description of the provisioning profile.
  - c. In the **Overallocation Limit (%)** field, type the overallocation limit percentage. Specify this value as an integer. The over allocation value identifies the limit on the virtual capacity that is allocated for a pool. A lower over allocation percentage is considered conservative because the virtual capacity is closer to the real capacity. A higher allocation percentage is considered aggressive because the virtual capacity is farther from the real capacity. For example, the over allocation limit that is set in the Aggressive profile is 300%. Pools that contain volumes or virtual disks that are created using this profile will be allocated up to three times the real capacity. If you use the Aggressive profile and the pool size is 15 GB, the virtual capacity allocated for the pool can be up to 45 GB.
4. In the **DS8000 Settings** section, specify the parameters for creating a thin provisioning profile for IBM System Storage DS8000.

#### Specify DS8000

Select the **Specify DS8000** check box to create a thin provisioning profile for System Storage DS8000.

#### Space-Efficiency Type

Specify the type of thin provisioning to use for System Storage DS8000. The options are:

### **Extent Space Efficient (ESE)**

An ESE logical volume has a set of virtual extents that are not associated with a repository. The logical volume is dynamically allocated and deallocated from the extent pool.

### **Track Space Efficient (TSE)**

A TSE logical volume has a set of virtual extents and is associated with a repository.

5. In the **SVC Settings** section, specify the parameters for creating a thin provisioning profile for IBM System Storage SAN Volume Controller.

#### **Specify SVC**

Select the **Specify SVC** check box to create a thin provisioning profile for SAN Volume Controller.

#### **Real Size (%)**

Enter the percentage of disk space to be allocated for the space-efficient copy of the selected VDisk. Setting the real space for a space-efficient copy protects the copy from going offline when its managed disk (MDisk) group runs out of space.

#### **Auto Expand**

Select a value to specify whether the thin provisioning profile volumes automatically expand their real capacities by allocating new extents from their MDisk group.

#### **Grain Size (KB)**

Select the grain size in KB for a space-efficient copy of the VDisk.

#### **Warning Level (%)**

Enter a threshold to send a warning message for the VDisk copy when the capacity that is used exceeds the specified percentage level.

6. In the **XIV Settings** section, specify the parameters for creating a thin provisioning profile for the XIV system.

#### **Specify XIV**

Select the **Specify XIV** check box to create a thin provisioning profile for the XIV system.

#### **Locking Behavior**

Select the **Locking Behavior** required if the thin provisioned pool runs out of space:

##### **Read Only**

If the thin provisioned pool has run out of capacity, all the volumes in the pool are locked and no write commands are allowed. This option is the default.

**No IO** If the thin provisioned pool has run out of capacity, no write or read commands are allowed.

7. Click **Apply** to save your changes. To exit without adding a thin provisioning profile, click **Cancel**.

## **Deleting a thin provisioning profile**

Learn how to delete a thin provisioning profile.

You can only delete user-defined thin provisioning profiles.

1. In the left pane, expand **Disk Manager > Profile Management > Thin Provisioning Profiles**. A list of the currently defined thin provisioning profiles is displayed in the right pane.

2. Select the thin provisioning profile you want to delete.
3. Click **Delete**.

## Modifying and viewing a thin provisioning profile

Learn how to modify and view a thin provisioning profile.

You can only modify user-defined thin provisioning profiles.

1. In the left pane, expand **Disk Manager > Profile Management > Thin Provisioning Profiles**. A list of the currently defined thin provisioning profiles is displayed in the right pane.
2. Click the magnifying glass icon next to a thin provisioning profile. The Thin Provisioning Profile Properties window displays.
3. View the details of the thin provisioning profile. You can modify any of the profile settings except for the **Profile Name**.
4. Click **Apply** to save your changes. To exit without modifying a thin provisioning profile, click **Cancel**.

## Managing storage subsystems

This topic contains information about storage-subsystem management tasks.


### Grouping storage subsystems

A storage-subsystem group is a user-defined set of storage subsystems. You might want to group storage subsystems in order to make it easier to run monitoring or reporting jobs.

**Note:** A storage subsystem can belong to only one storage-subsystem monitoring group, but can belong to multiple storage subsystem reporting groups.

### Creating a storage subsystem group:

This topic describes how to create a storage subsystem group.

1. In the left pane, expand **Disk Manager > Monitoring > Groups**.
2. Right-click **Storage Subsystem** and then click **Create Storage Subsystem Group**. In the right pane, the Create Storage Subsystem Group window opens.
3. In the **Description** field, type a short description of the group.
4. In the **Available** field, click the storage subsystems that you want to add to the group, and then click . The storage subsystems are added to the **Current Selections** field.
5. Click **File > Save**. A window opens requesting a name for the storage-subsystem group.
6. Type the name of the storage-subsystem group and click **OK**.

### Deleting a storage-subsystem group:

This topic describes how to delete a storage-subsystem group.

1. In the left pane, expand **Disk Manager > Monitoring > Groups > Storage Subsystems**. All existing storage-subsystem groups are listed.
2. Right-click the storage-subsystem group that you want to delete.
3. Click **Delete**. A window opens, asking if you are sure that you want to delete this group.
4. Click **Yes**.

## Editing a storage-subsystem group:

This topic describes how to edit a storage-subsystem group.

1. In the left pane, expand **Disk Manager > Monitoring > Groups > Storage Subsystem**. All existing storage-subsystem groups are listed.
2. Click the storage-subsystem group that you want to edit. The **Edit Storage Subsystem Group** window opens.
3. Make any necessary changes.
4. Click **File > Save**.

## Defining element managers

Element managers are programs that you can use to configure and manage your storage devices. An example of an element manager is DS8000 Storage Manager, which is the default element manager that is provided with IBM System Storage DS8000.

Element managers are managed in Tivoli Storage Productivity Center using the Configuration Utility. You can access the Configuration Utility from a storage system that is visible in the Storage Subsystems page of Tivoli Storage Productivity Center as shown in the following steps:

1. Expand **Disk Manager > Storage Subsystems** in the navigation tree. The Storage Subsystems page is displayed.
2. In the **Storage Subsystems** table, double-click a storage subsystem. The Details dialog page is displayed.
3. Click **Element Management**. The **Element Manager** tab in the Configuration Utility is displayed.

### Related concepts:

“Managing DS8000 storage systems using element managers” on page 598

Use the Configuration Utility in IBM Tivoli Storage Productivity Center to add and manage element managers for IBM System Storage DS8000. *Element managers* are applications that are used to configure and manage a storage device. You can add and manage connections to the element manager that is provided with DS8000, DS8000 Storage Manager, or you can specify other element managers for a device.

“Managing other storage systems using element managers” on page 605

Use the Configuration Utility in IBM Tivoli Storage Productivity Center to add and manage element managers for storage devices other than IBM System Storage DS8000. *Element managers* are applications that are used to configure and manage a storage device. You manage connections to the element managers that are provided with the storage device or you can specify other element managers for a device.

## Starting element managers

You can start element managers from a storage system that is visible on the Storage Subsystems page of Tivoli Storage Productivity Center. You can also start element managers from other points in Tivoli Storage Productivity Center as described in other areas of this documentation.

To start an element manager from the Storage Subsystem page, complete the following steps:

1. Expand **Disk Manager > Storage Subsystems** in the navigation tree. The Storage Subsystems page is displayed.
2. In the **Storage Subsystems** table, select a storage system.
3. Use one of the following methods to start the element manager:

- To start the default element manager for the storage system, select **Launch Element Manager**.
- To start another element manager for the storage system, right-click the system and click **Launch Element Manager**. Select the element manager that you want to start from the menu.

If the connection information for an element manager is valid, the element manager is started. If unsuccessful, an error message is displayed. The error messages that you receive describes one of the following problems:

- that the element manager at the specified IP address could not be started
- that the required connection information for the element manager has not been provided.
- that the element manager is not installed and configured correctly (for local element manager installations such as XIV Storage Manager)

If you are prompted to enter additional information such as an ID and password for the element manager, complete the following steps:

1. Click **OK** in the message dialog box. The Modify Element Manager dialog box is displayed.
2. Enter the required information for the element manager and click **Save**.

## Removing a storage subsystem

This topic describes how to remove a storage subsystem.

1. In the left pane, expand **Disk Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems page is displayed.
2. In the **Storage Subsystem** table, click a storage subsystem, and then click **Remove**. The Remove Storage Subsystem window opens.
3. Click **Yes**.

**Note:** You can also remove a storage subsystem using the following method: “Removing a storage system” on page 292.

## Removing a storage subsystem

This topic describes how to remove a storage subsystem.

1. In the left pane, expand **Disk Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems page is displayed.
2. In the **Storage Subsystem** table, click a storage subsystem, and then click **Remove**. The Remove Storage Subsystem window opens.
3. Click **Yes**.

**Note:** You can also remove a storage subsystem using the following method: “Removing a storage system” on page 292.

## Setting user-defined properties

This topic describes how to set properties for storage subsystems. You can define a name and up to three properties for each storage subsystem.

1. In the navigation tree, expand **Disk Manager** and click **Storage Subsystems**. In the content pane, the Storage Subsystems page is displayed.
2. In the **Storage Subsystem** table, double-click a storage subsystem. The Storage-Subsystem Details page opens.

| Storage Subsystems                           | DS8000-2107-1300361-IBM Details                   |
|----------------------------------------------|---------------------------------------------------|
| <b>Storage subsystem</b>                     | DS8000-2107-1300361-IBM                           |
| <b>Label</b>                                 | <input type="text"/>                              |
| <b>Status</b>                                | <input checked="" type="checkbox"/> Normal        |
| <b>Vendor</b>                                | IBM                                               |
| <b>Type</b>                                  | DS8000                                            |
| <b>Model</b>                                 | 921                                               |
| <b>Available Space (GB)</b>                  | 570.81                                            |
| <b>Consumed Space (GB)</b>                   | 5,394.07                                          |
| <b>Configured Real Space (GB)</b>            | 5,964.88                                          |
| <b>Available Real Space (GB)</b>             | 567.81                                            |
| <b>Serial number</b>                         | 1300361                                           |
| <b>Revision</b>                              | 5.4.30.253                                        |
| <b>User-defined property 1 (UDP1)</b>        | <input type="text"/>                              |
| <b>User-defined property 2 (UDP2)</b>        | <input type="text"/>                              |
| <b>User-defined property 3 (UDP3)</b>        | <input type="text"/>                              |
| <b>IP Address</b>                            | 9.47.97.252                                       |
| <b>Element manager</b>                       | https://9.47.97.252:8452/DS8000                   |
|                                              | <input type="button" value="Element Management"/> |
| <b>Replication Information</b>               |                                                   |
| The device is not configured for replication |                                                   |

Figure 44. Storage-Subsystem Details page

3. Set the user-defined properties:
  - a. In the **Label** field, type a name for the storage subsystem.
  - b. In the **User-defined properties** fields, type information that you want to associate with the storage subsystem.
4. Click **File > Save**.

### Viewing storage subsystem information

You can view information about the storage subsystems that are in your environment. This information includes the available and used space for the subsystem.

To view information about a storage subsystem:

1. In the navigation tree, expand **Disk Manager** and click **Storage Subsystems**. The Storage Subsystems page is displayed.  
The **Storage Subsystems** table contains the following information about each storage subsystem:

#### Subsystem

A system ID which is in the following format:  
*Type-Model-ID-Vendor*

- *Type* is the type of the storage system, for example, TotalStorage Enterprise Storage Server or SAN Volume Controller.
- *Model* is the model number, for example, 2105 or 2145.
- *ID* is the serial number or cluster name. Cluster names are used only for storage systems such as SAN Volume Controller and Storwize V7000.
- *Vendor* is the manufacturer of the storage system, for example, IBM.

**Label** A user-defined name or the system ID.

**Type** The type of system:

- DS4K
- DS5K
- DS6K
- DS8K
- EMC
- ESS
- Hitachi
- Storwize V7000
- Storwize V7000U
- SVC
- XIV
- Other

**Model** Model of the storage system.

**Status** Status of the system:

- Normal
- Critical
- Warning
- Unreachable

**Available Space (GB)**

Total storage capacity that is not formatted in volumes or virtual disks and assigned to hosts.

For the following storage systems, this value represents managed disk space that is not allocated to virtual disk space within the managed disk groups and managed disk space that is not defined to a managed disk group:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

**Tip:** For space-efficient volumes, this value is the amount of storage space requested for these volumes, not the actual allocated amount. This can result in discrepancies in the overall storage space reported for a storage system using space-efficient volumes. This also applies to other space calculations, such as the calculations for the storage system's Consumable Volume Space and FlashCopy Target Volume Space.

**Consumed Space (GB)**

Total storage capacity that is formatted in volumes or virtual disks and assigned to hosts.

For the following storage systems, this value represents all space that is allocated from MDisk groups to VDisks (space that is consumed by volumes):

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

**Configured Real Space (GB)**

Total amount of real available and formatted storage space in the storage pools that are associated with the storage system.

**Available Real Space (GB)**

Amount of real storage space available to allocate to volumes in the storage pools that are associated with the storage system.

**Tiers** Indicates the tiers that are used by the storage system. Each tier is delimited by a semi-colon (;). Possible values are: generic\_hdd and generic\_ssd.

Information about tiers is available only for the following storage systems:

- SAN Volume Controller 6.1 or later
- Storwize V7000
- Storwize V7000 Unified

The value N/A is displayed in tier-related columns for other storage systems.

**Tier Capacities (GB)**

Total storage capacity of each tier that is used by the storage system. The capacity of each tier is delimited by a semi-colon (;).

**Tier Free Capacities (GB)**

Free storage capacity (available space) of each tier that is used by the storage system. The free capacity of each tier is delimited by a semi-colon (;).

2. To view detailed information about a specific storage subsystem, click the

storage subsystem, and then click  . The Storage-Subsystem Details page opens.

The Storage-Subsystem Details page contains the following information:

**Storage Subsystem**

Storage system identifier.

**Label** User-defined name of the storage system. You can edit this field.

**Status** Status of the storage system:

- Normal
- Critical
- Warning
- Unreachable

**Vendor**

Manufacturer of the storage system.



**Type** Type of storage system:

- DS4K
- DS5K
- DS6K
- DS8K
- EMC
- ESS
- Hitachi
- Storwize V7000
- Storwize V7000U
- SVC
- XIV
- Other

**Model** Model of the storage system.

#### **Available Space**

Total storage capacity that is not formatted in volumes or virtual disks and assigned to hosts.

For the following storage systems, this value represents managed disk space that is not allocated to virtual disk space within the managed disk groups and managed disk space that is not defined to a managed disk group:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

**Tip:** For space-efficient volumes, this value is the amount of storage space requested for these volumes, not the actual allocated amount. This can result in discrepancies in the overall storage space reported for a storage system using space-efficient volumes. This also applies to other space calculations, such as the calculations for the storage system's Consumable Volume Space and FlashCopy Target Volume Space.

#### **Consumed Space**

Total storage capacity that is formatted in volumes or virtual disks and assigned to hosts.

For the following storage systems, this value represents all space that is allocated from MDisk groups to VDisks (space that is consumed by volumes):

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

#### **Configured Real Space (GB)**

Total amount of real available and formatted storage space in the storage pools that are associated with the storage system.

#### **Available Real Space (GB)**

Amount of real storage space available to allocate to volumes in the storage pools that are associated with the storage system.

**Tiers** Indicates the tier levels that are used by the storage system. Each tier is displayed as a row in a table. Information about tiers is available only for the following storage systems:

- SAN Volume Controller 6.1 or later
- Storwize V7000
- Storwize V7000 Unified

The value N/A is displayed in tier-related columns for other storage systems.

**Tier** Indicates a tier level that is used by the storage system. Possible values are: generic\_hdd and generic\_ssd.

**Tier Capacity (GB)**

Total storage capacity of a tier.

**Tier Free Capacities (GB)**

Free storage capacity (available space) of a tier.

**Serial Number**

Serial number.

**Revision**

Revision number.

**User-defined property 1 (UDP1)**

A user-defined property. You can edit this field.

**User-defined property 2 (UDP2)**

A user-defined property. You can edit this field.

**User-defined property 3 (UDP3)**

A user-defined property. You can edit this field.

**IP Address**

IP address of the element manager.

**Element Manager**

Identifies the element manager. This value can be a URL, the fully qualified name of the executable file that starts the element manager, or the short name for the XIV system.

**Element Management**

Opens the Configuration Utility. Use the Configuration Utility to modify the element manager information.

**Replication Information**

If the storage system is not configured for replication, one of the following messages is displayed:

- The device is not supported.
- The replication information for the device is unavailable.
- The device is not configured for replication.

If the storage system is configured for replication, the following information that is specific to the device type and connection type is displayed:

**IP Address/Domain Name**

Identifies the IP address or host name of the storage system. If the IP address was not defined, this field is blank.

**Port** Identifies the port number of the storage system. If the port

number was not defined, this field is blank. This field is not applicable for storage systems that use the HMC connection method.

**Username**

Identifies the user name to access the storage system. If the user name was not defined, this field is blank.

**Connection Type**

Identifies the connection method.

If you are viewing information about a Storwize V7000 Unified storage system, the following information is provided for Storwize V7000 File Module:

**Available File System Space**

The amount of available file system space for the file module.

**Consumed File System Space**

The amount of used file system space for the file module.

**Version**

The version of the file module.

## Managing virtual disks

This topic contains information about managing virtual disks.

### Adding a managed disk to a managed-disk group

This topic describes how to add a managed disk to a managed-disk group.

**Attention:** If you add a managed disk that contains existing data to a managed disk group, you will lose the data that it contains. The *image mode* is the only mode that will preserve this data.



1. In the left pane, expand **Disk Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems page is displayed.
2. In the **Storage Subsystems** table, click a SAN Volume Controller.
3. Click **Create Virtual Disk**. The Create Virtual Disk wizard starts and the "Define the virtual disks" page opens.
4. In the **Managed-disk group** field, select the managed-disk group to which you want to add a managed disk.
5. Click **Add Managed Disk**. The "Add managed disk" window opens.
6. Select the managed disk, and then click **OK**.

### Assigning host ports

This topic describes how to assign host ports to a virtual disk. You also can make host-port assignments when you create a virtual disk.

**Note:** When you assign a host port to a virtual disk, all host ports defined on the SAN Volume Controller model are assigned to the virtual disk. If you want to assign only one host port to a virtual disk, you must define the host with only one host port on the SAN Volume Controller.

1. In the left pane, expand **Disk Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems window is displayed.
2. In the **Storage Subsystems** table, click the SAN Volume Controller with which you want to work.
3. Click **Virtual Disks**. The Virtual Disks page is displayed.

4. Click either **Filter by MDisk group** or **Filter by host port**, and then select the specific filtering condition.
5. Click **Go**. A list of virtual disks associated with the selected SAN Volume Controller is displayed.
6. In the **Virtual Disks** table, select the virtual disk, and then click **Assign Host Ports**. The Assign Host Ports wizard starts. A virtual disk that is marked with an asterisk (\*) indicates that the virtual disk is a space-efficient volume. The \* is not part of the actual virtual disk name.
7. On the "Assign the virtual disk to host ports" page, assign or unassign host ports:
  - a. Optional: In the **Available ports** field, select the ports that you want to assign to the virtual disks, and then click .
  - b. Optional: In the **Assigned ports** field, select any ports that you want to unassign from the virtual disks, and then click .
  - c. Click **Next**.
8. (This page only is displayed if Fabric Manager is installed, and you have Fabric administrator authority.) On the "Zone actions" page, perform one of the following zone actions:
  - Create a new zone**  
Click **Create new Zone**. In the Zone name field, type the name of the zone. The ports listed in the field below are added to the new zone.
  - Update an existing zone**  
Click **Update existing zone**, and then select a zone. The ports listed in the field below are added to the selected zone.
  - Leave the zone configuration as it is**  
Click **Do neither** to avoid performing any fabric-control tasks. The ports listed in the field below cannot communicate with the subsystem on this fabric.
9. Click **Next**. The "Review settings" page opens.
10. Review the settings that you made earlier in the wizard, and then click **Finish**.

## Creating virtual disks

This topic describes how to create virtual disks.

1. Expand the **Disk Manager** node in the navigation tree and click **Storage Subsystems**. In the right pane, the Storage Subsystems page is displayed.
2. In the **Storage Subsystem** table, click a IBM System Storage SAN Volume Controller, and then click **Create Virtual Disk**. The Create Virtual Disk wizard is started and the Define the virtual disks page is opened.

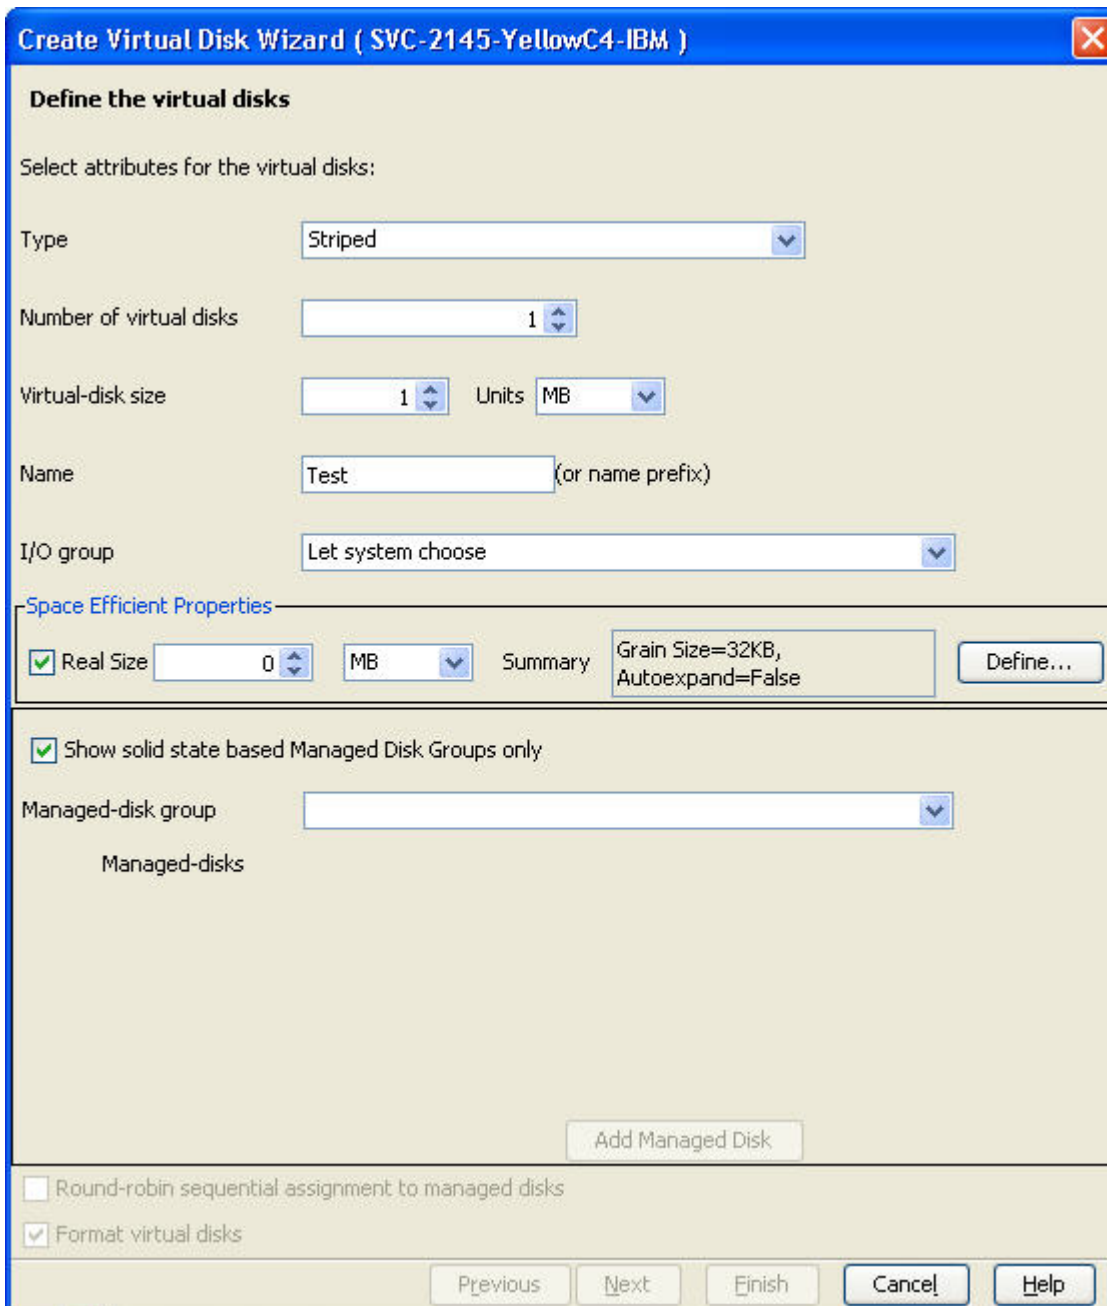



Figure 45. Create Virtual Disk wizard: "Define the virtual disks" page

3. On the Define the virtual disks page, specify the attributes for the virtual disks:
  - a. In the **Type** field, select one of the following values:
    - Striped
    - Sequential
  - b. In the **Number of virtual disks** field, select the number of virtual disks that you want to create. Each available input/output (I/O) group can contain 1024 virtual disks.
  - c. In the **Virtual-disk size** field, select the size of the virtual disks to be created.

- d. In the **Units** field, select one of the following units of measurement for the size of the virtual disks:
- Bytes (byte)
  - GB (gigabyte)
  - KB (kilobyte)
  - MB (megabyte)
  - PB (petabyte)
  - TB (terabyte)
- e. In the **Name** field, type a string to identify the virtual disk.
- For SAN Volume Controller 4.x and 5.x, the following naming conventions apply:
    - The maximum length of the name is 14 characters. The number of virtual disks also determines the maximum length of the name. For example, if the number of virtual disks is 10 - 99, the maximum length of the name is 13 characters. If the number of virtual disks is 100 - 999, the maximum length of the name is 12 characters, and so on.
    - Permissible characters include uppercase letters (A-Z), lowercase letters (a-z), numerals (0-9), hyphens (-), and underscores (\_).
  - For SAN Volume Controller 6.1 and Storwize V7000, the following naming conventions apply:
    - The maximum length of a single virtual disk is 63 characters. The number of virtual disks also determines the maximum length of the name. For example, if the number of virtual disks is 2 - 9, the maximum length of the name is 62 characters. If the number of virtual disks is 10 - 99, the maximum length of the name is 61 characters. If the number of virtual disks is 100 - 999, the maximum length of the name is 60 characters, and so on.
    - Permissible characters include uppercase letters (A-Z), lowercase letters (a-z), numerals (0-9), hyphens (-), underscores (\_), spaces(), and periods(.
  - Names must begin with a letter or an underscore. The first character cannot be a number, hyphen (-), space( ), or period(.
  - Names cannot begin with text or abbreviations commonly used to specify the object type, for example, the name cannot begin with VDisk, VDISK, vdisk, or virtual disk.
  - Names cannot end with a space.
  - Object names must be unique within an object type. For example, it is acceptable to have a vdisk named "XYZ" and an mdisk named "XYZ", but not acceptable to have two vdisks named "XYZ".
- If you leave this field blank, a name is generated automatically. If you create multiple virtual disks, the name is used as a prefix.
- f. In the **I/O group** field, select the I/O group. By default, **Let the system choose** is selected. Other options are the actual I/O groups on the cluster.
- g. In the **Space-efficient properties** pane, use the following options to specify the space-efficient properties of the virtual disk. This pane is disabled if the SAN Volume Controller does not support the creation of space-efficient virtual disks:
- 1) Space efficient check box (precedes the **Real size** field). Select this check box to specify that the virtual disk will be space efficient. Selecting this check box activates the other fields in this section.

- 2) **Real size** Specify the actual allocated size of the virtual disk at creation time. This size is the initial allocated disk size that can be increased as needed by the SAN Volume Controller when the applications using the VDisk need more storage. The real size of the disk cannot exceed the managed free space that is currently available. Use the dropdown list to specify the size of the allocation unit.
  - 3) **Summary** This field displays a summary of the space-efficient properties for this virtual disk.
  - 4) **Define** Click this button to open the Define Space Efficient Properties window.
- h. Select the **Show solid state based Managed Disk Groups only** check box if you want only managed disk groups that are on solid-state drives to be displayed in the **Managed-disk group** field. Solid-state drives are available in SAN Volume Controller 5.1 and later.
  - i. In the **Managed-disk group** field, select the managed disk group. The managed disk groups are listed using the following convention:  
MDisk-group name (*N*)  
  
where *N* is the available capacity. After you select the managed disk group, the managed disks in the managed disk group are displayed.  
To add a managed disk to the managed disk group, click **Add Managed Disk**.
  - j. This check box is available only for sequential virtual disks. If you are creating more than one sequential virtual disk, select the **Round robin sequential assignment to managed disks** check box to ensure that the virtual disks are assigned to the managed disks in a round-robin manner. If selected, the number of virtual disks must equal the number of managed disks selected in the managed disk table.  
  
Clear the **Round robin sequential assignment to managed disks** check box to ensure that all virtual disks are created on a single managed disk. If the check box is cleared, there must be *only* one managed disk selected in the managed-disk table.
  - k. Select the **Format virtual disks** check box to ensure that the virtual disks are formatted (written with zeros) after they are created.
  - l. Click **Next**.
4. On the Assign the virtual disk to host ports page, assign host ports:
    - a. Optional: In the **Available ports** field, select the ports that you want to assign to the virtual disks, and then click  .
    - b. Click **Next**.
  5. On the Zone actions page, perform one of the following zone actions. The Zone actions page is displayed only if Fabric Manager is installed, and you have Fabric administrator authority.

#### Create a new zone

Click **Create new Zone**. In the **Zone name** field, type the name of the zone. The ports that are listed are added to the new zone.

#### Update an existing zone

Click **Update existing zone**, and then select a zone. The ports that are listed are added to the selected zone.

### Leave the zone configuration as it is

Click **Do neither** to avoid performing any fabric-control tasks. The ports that are listed are not able to communicate with the subsystem on this fabric.

6. Click **Next**.
7. On the Summary page, review the selections that you made earlier in the wizard.
8. Click **Finish** to create the virtual disks.

### Deleting a virtual disk

This topic describes how to delete a virtual disk.

1. In the left pane, expand **Disk Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems window is displayed.
2. In the **Storage Subsystems** table, click a SAN Volume Controller, and then click **Virtual Disks**. The Virtual Disks page opens.
3. Click either **Filter by MDisk group** or **Filter by host ports**, and then select the specific filtering condition.
4. Click **Go**. A list of virtual disks associated with the selected SAN Volume Controller is displayed.

| Name     | I/O Group | Status   | MDisk Group Name | Capacity (MB) | Type       | FlashCopy Pair Name | Metro Mirror Name |
|----------|-----------|----------|------------------|---------------|------------|---------------------|-------------------|
| c20_load | io_grp0   | Degraded | mdiskgrp4        | 1,477         | Striped    |                     |                   |
| mdm22    | io_grp0   | Degraded | mdiskgrp4        | 0.95          | Striped    |                     |                   |
| my       | io_grp0   | Degraded | mdiskgrp4        | 0.95          | Striped    |                     |                   |
| mydisk   | io_grp0   | Degraded | mdiskgrp4        | 0.95          | Striped    | fcmap2              |                   |
| rhvdisk  | io_grp0   | Degraded | mdiskgrp4        | 9.54          | Striped    |                     |                   |
| test123  | io_grp0   | Degraded | mdiskgrp4        | 0.95          | Striped    |                     |                   |
| vdisk35  | io_grp0   | Degraded | mdiskgrp4        | 64            | Striped    |                     |                   |
| vdisk36  | io_grp0   | Degraded | mdiskgrp4        | 64            | Striped    |                     |                   |
| vdisk37  | io_grp0   | Degraded | mdiskgrp4        | 64            | Sequential |                     |                   |
| vdisk38  | io_grp0   | Degraded | mdiskgrp4        | 953.67        | Image      |                     |                   |
| vdisk48  | io_grp0   | Degraded | mdiskgrp4        | .1            | Striped    |                     |                   |

Figure 46. Virtual Disks page

5. In the **VDisk** table, click the virtual disk that you want to delete, and then click **Delete**.
6. When prompted to confirm that you want to delete the virtual disk, click **Yes**.

### Viewing virtual-disk information

This topic describes how to view information about the virtual disks in your environment.

1. Expand the **Disk Manager** node in the navigation tree and click **Storage Subsystems**. In the right pane, the Storage Subsystems page is displayed.
2. In the **Storage Subsystems** table, click a SAN Volume Controller, and then click **Virtual Disks**. The Virtual Disks page opens.



3. Click either **Filter by MDisk group** or **Filter by host ports**, and then select the specific filtering condition.
4. Click **Go**. A list of virtual disks associated with the selected SAN Volume Controllers is displayed. The **VDisk** table contains the following information about each virtual disk:

**Name** Name of the virtual disk or mirror copy.

**I/O Group**

Name of the I/O group that is associated with the virtual disk or mirror copy.

**Status** Status of the virtual disk or mirror copy:

- Offline
- Online
- Degraded

**Copy Type**

Indicates the type of the virtual disk. Valid values for this column include the following:

- **Standard:** Indicates standard VDIs (from pre-SAN Volume Controller 4.3 systems)
- **Primary:** Indicates VDIs from SAN Volume Controller 4.3 or newer systems.
- **Copy:** Indicates a mirror copy (not the primary copy) of a VDI.

**Note:** You can assign virtual disks to hosts if those disks have a **Copy Type** of **Standard** or **Primary**.

**Copy ID**

Indicates the ID of a mirror copy for a virtual disk. Depending on the configuration history, a single mirror copy might have an ID of either 0 or 1.

**Note:** The value N/A appears in this column for virtual disks.

**MDisk Group Name**

Name of the managed-disk groups to which the virtual disk or mirror copy belongs.

**Capacity (MB)**

Capacity of the virtual disk or mirror copy in MB. This value is rounded to 0.1 MB.

**Type** Type of virtual disk or mirror copy:


- Image
- Sequential
- Striped

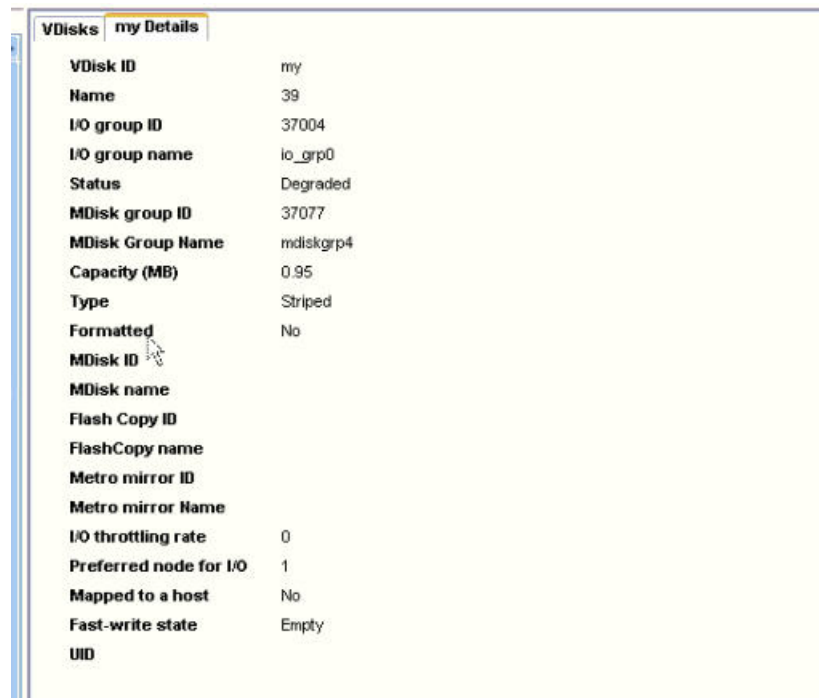
**FlashCopy Name**

Name of the FlashCopy mapping to which the virtual disk belongs.

**MM Name**

Name of the Metro Mirror relationship to which the virtual disk belongs.

- To view detailed information about a specific virtual disk, click the virtual disk, and then click . The Virtual-Disk Details page opens.



| VDisks   my Details           |           |
|-------------------------------|-----------|
| <b>VDisk ID</b>               | my        |
| <b>Name</b>                   | 39        |
| <b>I/O group ID</b>           | 37004     |
| <b>I/O group name</b>         | io_grp0   |
| <b>Status</b>                 | Degraded  |
| <b>MDisk group ID</b>         | 37077     |
| <b>MDisk Group Name</b>       | mdiskgrp4 |
| <b>Capacity (MB)</b>          | 0.95      |
| <b>Type</b>                   | Striped   |
| <b>Formatted</b>              | No        |
| <b>MDisk ID</b>               |           |
| <b>MDisk name</b>             |           |
| <b>Flash Copy ID</b>          |           |
| <b>FlashCopy name</b>         |           |
| <b>Metro mirror ID</b>        |           |
| <b>Metro mirror Name</b>      |           |
| <b>I/O throttling rate</b>    | 0         |
| <b>Preferred node for I/O</b> | 1         |
| <b>Mapped to a host</b>       | No        |
| <b>Fast-write state</b>       | Empty     |
| <b>UID</b>                    |           |

Figure 47. Virtual-Disk Details page

The Virtual-Disk Details page contains the following information:

**VDisk ID**

ID of the virtual disk.

**Name** Name of the virtual disk. A virtual disk that is marked with an asterisk (\*) indicates that the virtual disk is a space-efficient volume. The \* is not part of the actual virtual disk name.

**I/O group ID**

ID of the I/O group.

**I/O group name**

Name of the I/O group.

**Status** Status of a virtual disk:

- Offline
- Online
- Degraded

**MDisk group name**

Name of the managed-disk group.

**Capacity (MB)**

Capacity of a virtual disk in MB. This value is rounded to 0.1 MB.

**Easy Tier**

Indicates the enablement of tiering. Possible values are: auto, on, off.

Information about virtual disk tiers is available only for:

- SAN Volume Controller 6.1
- Storwize V7000
- Storwize V7000 Unified

The fields related to tiers are not displayed for other storage systems.

**Easy Tier Status**

Indicates the status of tiering. Possible values are: active, inactive.

**Tiers** Indicates the tier levels used by the virtual disk. Each tier is displayed as a row in a table. Possible values are: generic\_hdd, generic\_ssd.

**Tier Capacity**

Total storage capacity of a tier.

**Type** Type of a virtual disk:

- Image
- Sequential
- Striped

**Formatted**

Indicates if a virtual disk is formatted.

**MDisk ID**

IDs of the MDisks that a virtual disk is using.

**MDisk name**

Names of the MDisks that a virtual disk is using.

**FlashCopy**

Indicates if the virtual disk is in a FlashCopy (FC) relationship and whether it is an FC source or FC target. Virtual disks that are not in an FC relationship are displayed with a None value whether or not the virtual disk is on a system that supports FC.

**Note:**

- This value is available for virtual disks of the following systems only: IBM TotalStorage Enterprise Storage Server, IBM System Storage DS6000, IBM System Storage DS8000, and SAN Volume Controller.
- TotalStorage Enterprise Storage Server systems must have at least the following microcode levels: ess800mincodelevel = 2.4.3.56, essf20mincodelevel = 2.3.3.89. IBM Tivoli Storage Productivity Center does not report FC information if the TotalStorage Enterprise Storage Server systems do not meet this requirement.

**FlashCopy ID**

ID of the FlashCopy mapping to which the virtual disk belongs.

**FlashCopy name**

Name of the FlashCopy mapping to which the virtual disk belongs.

**Metro mirror ID**

ID of the Metro Mirror relationship to which the virtual disk belongs.

**Metro mirror name**

Name of the Metro Mirror relationship to which the virtual disk belongs.

**I/O throttling rate**

Maximum rate at which an I/O transaction is accepted for a virtual disk.

**Preferred node for I/O**

IP address of the node to be used for I/O operations.

**Mapped to a host**

Yes or No.

**Fast-write state**

Yes or No. This value indicates whether the virtual disk has data in the cache.

**UID** Unique ID.

## Managing volumes

This topic contains information about managing volumes.

### Assigning host ports

This topic describes how to assign host ports to volumes. You also can make host-port assignments when you create volumes.

1. In the left pane, expand **Disk Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems window is displayed.
2. In the **Storage Subsystems** table, click the storage subsystem with which you want to work.
3. Click **Volumes**. The Volumes page is displayed.

**Volumes** FSVT\_1000 Details

Create Assign Host Ports Delete

A volume job is in progress. To view the job status, click Monitoring --> Jobs.

Volume view:

Filter by extent pool 2107.1301901-vs0 Go


Filter by host port Go

Volumes (2107.1301901):

| Volume           | Subsystem                      | Subsystem Name      | LCU       | Type      | Size        | Unit |
|------------------|--------------------------------|---------------------|-----------|-----------|-------------|------|
| Camera_1809      | DS8000-2107-1301901-IBM        | 2107.1301901        | 24        | FB        | 5 GB        |      |
| Camera_180A      | DS8000-2107-1301901-IBM        | 2107.1301901        | 24        | FB        | 5 GB        |      |
| Camera_180B      | DS8000-2107-1301901-IBM        | 2107.1301901        | 24        | FB        | 5 GB        |      |
| Cuda_3606        | DS8000-2107-1301901-IBM        | 2107.1301901        | 54        | FB        | 5 GB        |      |
| Cuda_3607        | DS8000-2107-1301901-IBM        | 2107.1301901        | 54        | FB        | 5 GB        |      |
| Cuda_3608        | DS8000-2107-1301901-IBM        | 2107.1301901        | 54        | FB        | 5 GB        |      |
| <b>FSVT_1000</b> | <b>DS8000-2107-1301901-IBM</b> | <b>2107.1301901</b> | <b>16</b> | <b>FB</b> | <b>5 GB</b> |      |
| FSVT_1001        | DS8000-2107-1301901-IBM        | 2107.1301901        | 16        | FB        | 5 GB        |      |
| FSVT_1002        | DS8000-2107-1301901-IBM        | 2107.1301901        | 16        | FB        | 5 GB        |      |
| FSVT_1003        | DS8000-2107-1301901-IBM        | 2107.1301901        | 16        | FB        | 5 GB        |      |
| FSVT_1004        | DS8000-2107-1301901-IBM        | 2107.1301901        | 16        | FB        | 5 GB        |      |
| FSVT_1005        | DS8000-2107-1301901-IBM        | 2107.1301901        | 16        | FB        | 5 GB        |      |
| FSVT_1006        | DS8000-2107-1301901-IBM        | 2107.1301901        | 16        | FB        | 5 GB        |      |
| FSVT_1007        | DS8000-2107-1301901-IBM        | 2107.1301901        | 16        | FB        | 5 GB        |      |

Figure 48. Volumes page

4. Select a filtering condition, and then click **Go**. A list of volumes is displayed.
5. In the **Volumes** table, select the volume, and then click **Assign Host Ports**. The Assign Host Ports wizard starts.
6. On the "Assign the volume to host ports" page, assign or unassign host ports:
  - a. Optional: In the **Available ports** field, select the ports that you want to assign to the virtual disks, and then click **>>**.

- b. Optional: In the **Assigned ports** field, select any ports that you want to unassign from the virtual disks, and then click . To avoid lost or inconsistent data, make sure that the volume is unmounted before unassigning.  
**Note:** if a host (agent) with multiple host ports is defined to a storage subsystem, assigning a volume to the host ports will create a single host definition with all of those ports.  
If a host (agent) with multiple host ports is **not** defined to a storage subsystem on the same fabric, assigning a volume to the host ports will create a separate host definition for each port (multiple host definitions, each with a single port).
  - c. Click **Next**.
7. On the "Assign the volume to subsystem ports" page, assign the volumes to specific subsystem ports.
    - a. Optional: In the **Available ports** list, select subsystem ports and click >> to assign them to the volumes. The ports are identified using the following convention: <Worldwide port name> (<host name>).
    - b. Click **Next**.
  8. (This page only is displayed if Fabric Manager is installed, and you have Fabric administrator authority.) On the "Zone actions" page, perform one of the following zone actions:
    - Create a new zone**  
Click **Create new Zone**. In the Zone name field, type the name of the zone. The ports listed in the field below are added to the new zone.
    - Update an existing zone**  
Click **Update existing zone**, and then select a zone. The ports listed in the field below are added to the selected zone.
    - Leave the zone configuration as it is**  
Click **Do neither** to avoid performing any fabric-control tasks. The ports listed in the field below will not be able to communicate with the subsystem on this fabric.
  9. Click **Next**. The "Review settings" page opens.
  10. Review the selections that you made earlier in the wizard, and then click **Finish**.

## Creating volumes and assigning volumes to a host



This task describes how to create volumes and assign the volumes to a host storage system. This task also describes how to perform zone configuration tasks if Fabric Manager is installed, and you have Fabric administrator authority

The performance of a volume creation job and the general performance of a CIM object manager (CIMOM) might decrease if you create many volumes on a system at the same time that a performance monitor is collecting data about that system. To avoid this issue, create volumes on a system when a performance monitoring job is not running or is not scheduled to run on the target system.

The following steps describe how to create volumes, assign the volumes to a storage system. Some fields described in these steps are dependent on the type of storage system that you have selected.

1. Expand **Disk Manager** and click **Storage Subsystems**.

2. In the **Storage Subsystem** table, select a storage system, and click **Create Volume**. The Create Volume wizard is started.
3. On the Define the volumes page, specify the attributes for the volumes:
  - a. In the **RAID level** field, select one of the following Redundant Array of Independent Disks (RAID) levels, based on the type of storage system:
    - IBM System Storage DS6000 and IBM System Storage DS8000: RAID 5, RAID 6 (DS8000 only), RAID 10, or Mixed. RAID 5 is the default setting.
    - IBM TotalStorage Enterprise Storage Server: RAID 5 or RAID 10. RAID 5 is the default setting.
    - IBM System Storage FAStT/DS4000 and IBM System Storage DS5000 Storage Manager: RAID 0, RAID 1, RAID 3, or RAID 5. RAID 0 is the default setting.
  - b. In the **Encryption group** field, select an encryption group if you want the new volumes to be created on extent pools that use full disk encryption. If an encryption group is selected, only pools that belong to the selected encryption group are displayed in the **Available extent pools** field. If there are no encryption groups available, **No selection** is displayed in the **Encryption group** field.
  - c. Select the **Solid state based pools only** check box if you want the new volumes to be created on extent pools that are on solid-state drives. If this check box is selected, only pools on solid-state drives are displayed in the **Available extent pools** field. If there are no solid-state drives available, the **Solid state based pools only** check box is disabled.
  - d. In the **Available extent pools** field, select an extent pool from the list of extent pools with the specified RAID level. These extent pools are identified using the following convention: *Device-specific extent pool name (free space GB)*. This field is displayed if the system is DS6000 or DS8000. Only extent pools with fixed-block (FB) tracks are listed.
  - e. In the **Available arrays** field, select an array from the list of arrays with the specified RAID level. These arrays are identified using the following convention: *Device-specific array name (free space GB)*. This field is displayed if the system is neither DS6000 nor DS8000.
  - f. In the **Number of volumes** field, select the number of volumes that you want to create. The default is 1. For an TotalStorage Enterprise Storage Server system, the maximum is 255.
  - g. In the **Volume size** field, select the size of the volumes to be created. The maximum and minimum size depends on the type of storage system.
  - h. In the **Units** field, select one of the following units of measurement for the size of the volumes:
    - KB
    - MB
    - GB
  - i. In the **Name** field, type a name for the volume. If you are creating multiple volumes, this string is used as a prefix. The string cannot be longer than 15 characters. This field is disabled for TotalStorage Enterprise Storage Server systems. TotalStorage Enterprise Storage Server volume names are determined by TotalStorage Enterprise Storage Server.
  - j. Click **Next**.
4. On the Assign the volume to host ports page, assign the volume to specific host ports:

- a. In the **Available ports** list, select host ports and click  to assign the ports to the volumes. The ports are identified using the following convention: *Worldwide port name (host name)*.
  - b. Click **Next**.
5. On the Restrict access to subsystem ports page, you can restrict the host access to specific system ports:
    - a. Optional: In the **Unassigned ports** list, select system ports and click  to remove the ports from the list of ports assigned to the volume. The ports are identified using the following convention: *Worldwide port name (host name)*. To avoid lost or inconsistent data, make sure that the volume is unmounted before you unassign ports.
    - b. Click **Next**.
  6. On the Zone actions page, perform one of the following zone actions. This page is displayed only if Fabric Manager is installed, and you have Fabric administrator authority.

Use a Zone Action if you want to apply updates to the active zone set. This action applies changes to the inactive zone set that has the same name as the active zone set and then the zone set is activated. Before using this option, ensure that the zones and zone members in the active zone set matches the zones and zone members in the inactive zone set with the same name.

#### Create a new zone

Click **Create new Zone**. In the Zone name field, type the name of the zone. The ports that are listed are added to the new zone.

#### Update an existing zone

Click **Update existing zone**, and then select a zone. The ports that are listed are added to the selected zone.

#### Leave the zone configuration as it is

Click **Do neither** to avoid performing any fabric-control tasks. The ports listed that are listed are not able to communicate with the subsystem on this fabric.

Click **Next**.

7. On the Summary page, review the selections that you made earlier in the wizard.
8. Click **Finish**.

## Deleting volumes

This topic describes how to delete volumes.

Note that volumes that are part of a replication session cannot be deleted.

1. In the left pane, expand **Disk Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems page is displayed.
2. In the **Storage Subsystems** table, click a storage subsystem, and then click **Volumes**. The Volumes page is displayed.
3. On the Volumes page, select the specific filtering condition, and then click **Go**. A list of volumes on the selected storage subsystem is displayed.
4. In the **Volumes** table, select the volume that you want to delete, and then click **Delete**. You can select one volume at a time only. You cannot delete TotalStorage Enterprise Storage Server volumes from within IBM Tivoli Storage Productivity Center.

5. When prompted to confirm that you want to delete the volume, click **Yes**. The volume is deleted automatically when you click Yes. You can reuse the volume name immediately after the volume is deleted.

**Note:** For SAN Volume Controller (SVC), if you choose to delete a volume using the SVC console rather than the IBM Tivoli Storage Productivity Center user interface, that volume is retained for 14 days before being deleted. You cannot reuse the volume name during that time.

## Viewing volume information

This topic describes how to view information about the volumes in your environment.

**Note:** If a probe is run against a CIM object manager (CIMOM) that is not available, the volumes defined on the storage system are not displayed in the IBM Tivoli Storage Productivity Center graphical user interface (GUI). Check that probes run successfully and consider configuring an alert to ensure that you are notified when a scheduled probe fails.

1. Expand the **Disk Manager** node in the navigation tree and click **Storage Subsystems**. In the right pane, the Storage Subsystems page is displayed.
2. In the **Storage Subsystems** table, click a storage system, and then click **Volumes**. The Volumes page opens.
3. Select the specific filtering condition, and then click **Go**. A list of volumes associated with the selected storage system is displayed. The **Volume** table contains the following information about each volume:

### Volume

Volume ID.

### Subsystem

A system ID in the following format:

*Type-Model-ID-Vendor*

- *Type* is the type of the storage system, for example, TotalStorage Enterprise Storage Server or SAN Volume Controller.
- *Model* is the model number, for example, 2105 or 2145.
- *ID* is the serial number or cluster name. Cluster names only are used for SAN Volume Controllers.
- *Vendor* is the manufacturer of the storage system, for example, IBM.

### Subsystem Label

A user-defined name or the system ID.

**LCU** Logical control unit (LCU), which can have one of the following values:

- DS6000, DS8000, and TotalStorage Enterprise Storage Server -- logical system (LSS) ID
- SAN Volume Controller --1

**Type** Type of volume, either fixed-block (FB) or count-key data (CKD).

### Volume Space

Size of the volume.

**Unit** GB for fixed-block (FB) volumes and cyl for count-key data (CKD) volumes.



4. To view detailed information about a specific volume, click the volume, and



. The Volume Details page is opened. The Volume Details page contains the following information:

**Volume**

Volume ID.

**Volume size**

Volume size.

**Volume real space**

Physical allocated space of the volume. For normal volumes, this value is equal to the volume space. For space efficient/thin provisioned volumes, this value is equal to the real space allocated when data is written to the volume. If IBM Tivoli Storage Productivity Center detects that a volume is thin provisioned but cannot retrieve the Volume real space value, N/A is displayed.

**Array** Array on which the volume is created.

**RAID level**

Redundant Array of Independent Disks (RAID) level of the volume.

**FlashCopy**

Indicates if the volume is in a FlashCopy (FC) relationship and whether it is an FC source or FC target. Volumes that are not in an FC relationship are displayed with a **None** value whether or not the volume is on a system that supports FC.

**Note:**

- This value is available for volumes of the following systems only: IBM TotalStorage Enterprise Storage Server (TotalStorage Enterprise Storage Server), IBM System Storage DS6000, IBM System Storage DS8000, and IBM System Storage SAN Volume Controller.
- TotalStorage Enterprise Storage Server systems must have at least the following microcode levels: `ess800mincodelevel = 2.4.3.56`, `essf20mincodelevel = 2.3.3.89`. IBM Tivoli Storage Productivity Center does not report FC information if the TotalStorage Enterprise Storage Server systems do not meet this requirement.

**Subsystem**

A system ID in the following format:

*Type-Model-ID-Vendor*

where:

- *Type* is the type of the storage system, for example, TotalStorage Enterprise Storage Server or SAN Volume Controller.
- *Model* is the model number, for example, 2105 or 2145.
- *ID* is the serial number or cluster name. Cluster names only are used in for SAN Volume Controllers.
- *Vendor* is the manufacturer of the storage system, for example, IBM.

**Subsystem Label**

This value can be either a user-defined name or the system ID.

**Is Encryptable**

Indicates **Yes** when the volume is encryptable and **No** when the volume is not encryptable.

**Is Encrypted**

Indicates **Yes** when the volume is encrypted and **No** when the volume is not encrypted.

**Assigned host ports**

Host ports to which the volume is assigned.

**Subsystem ports**

System ports to which the volume is assigned.

---

## Managing fabrics

This section describes operations available for directly manipulating the zone configuration in a fabric.

Zone control functions are also integrated in the Storage Provisioning function (see “Provisioning planner” on page 111) and the SAN Planner functions (see “SAN Planner overview” on page 417).

**Note:**

- After you click **Change** to begin your zone configuration changes and before you click **Update and Activate** or **Update Only** to conclude your changes, it is possible that other users might have made zoning changes to the same fabric. This situation might have happened through a switch element manager or switch CLI, as some switch vendors and models lack zone control session locking across the fabric. This lack of zone control session locking might allow overwrite conditions of the changes to the active zoneset for the fabric and to the set of inactive zonesets.
- If the zoning on a fabric has an active configuration and you use a switch element manager to delete the inactive copy of the active zoneset, Tivoli Storage Productivity Center still displays the inactive copy of the active zoneset on the Fabric Manager > Fabrics > Zone Configuration panels. Tivoli Storage Productivity Center displays the inactive copy even though it does not exist in the fabric zone configuration so that you can edit the active zoneset.

The contents of this inactive copy of zoneset displayed is identical to contents of the active copy of the zoneset. If you use Tivoli Storage Productivity Center to make any subsequent zoning changes for the fabric, the inactive copy of the active zoneset is created in the fabric zone configuration.

## Working with zones

This topic describes how to create, change, duplicate, activate or deactivate, and delete zones.

**Creating zones**

When you add zones, you assign assets to the zone and specify the zone set or sets to which the zone belongs.

To create a zone:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zones** area, click **Add**.
5. In the Zone Configuration window enter the name of the zone and, optionally, a description. The zone name must be unique in the fabric. Click **Next**.



6. In the next Zone Configuration window, select the assets (switches, devices, aliases) that you want to include in the zone, move them from the **Available Assets** column to the **Selected Assets** column, and click **Next**.  
If the fabric is set on Interop Mode, you can select only the port WWN number and the device port number from the **Available Assets** column.
7. In the next Zone Configuration window, select one or more zone sets, move them from the **Available Zone Sets** column to the **Selected Zone Sets** column, and click **Finish**.


## Adding storage to a server and setting up a zone

To add storage to a server and set up a zone, do the following:

1. **Add the CIMOM**
  - a. Expand **Administrative Services** → **Data Sources** → **CIMOM Agents** .
  - b. Click **Add CIMOM** in the right pane.
  - c. In the **Add CIMOM** window, enter the required values and click **Save**.
2. **Run a CIMOM discovery**
  - a. Expand **Administrative Services** → **Discovery**.
  - b. Right-click **CIMOM** and select **Run Now**.
  - c. Right-click **CIMOM** and select **Refresh Job List** for the submitted job to appear in the list of discovery jobs. The date and time of the job is displayed in the left pane. To see the latest status of the job, right-click **CIMOM** and select **Update Job Status**.
3. **Run a probe**

When the discovery has completed successfully, perform the following steps:

  - a. Expand **IBM Tivoli Storage Productivity Center** → **Monitoring** .
  - b. Right-click **Probes** and select **Create Probe**.
  - c. Make your selections and click **Save** to submit the job.
  - d. Right-click **Probes** and select **Refresh Job List** to display the job in the list of probes. The date and time of the job will be displayed in the left pane. To see the latest status of the job, right-click **Probes** and select **Update Job Status**.
4. In the left pane, expand **Data Manager** and click **Storage Subsystems**. In the right pane, the Storage Subsystems window opens.
5. In the **Storage Subsystems** table, click the storage subsystem with which you want to work.
6. Click **Volumes**.
7. In the Volumes page select a filtering condition, and click **Go**.
8. In the **Volumes** table, select the volume, and then click **Assign Host Ports**. The Assign Host Ports wizard starts.
9. On the "Assign the volume to host ports" page, you can assign and unassign host ports:
  - a. In the **Available ports** field, select the ports that you want to assign to the virtual disks, and then click  .
  - b. In the **Assigned ports** field, select any ports that you want to unassign from the virtual disks, and then click  .
  - c. Click **Next**.




10. On the "Assign the volume to subsystem ports" page, assign the volumes to specific subsystem ports.
  - a. Optional: In the **Available ports** list, select subsystem ports and click  to assign them to the volumes. The ports are identified using the following convention: <Worldwide port name> (<host name>).
  - b. Click **Next**.
11. "Zone actions" page is displayed only if Fabric Manager is installed, and you have Fabric administrator authority. Perform one of the following zone actions:
  - Create a new zone**  
Click **Create new Zone**. In the Zone name field, type the name of the zone. The ports listed in the field below are added to the new zone.
  - Update an existing zone**  
Click **Update existing zone**, and then select a zone. The ports listed in the field below are added to the selected zone.
  - Leave the zone configuration as it is**  
Click **Do neither** to avoid performing any fabric-control tasks. The ports listed in the field below will not be able to communicate with the subsystem on this fabric.


Click **Next**.
12. In the "Review settings" page, review your selections and click **Finish**.

## Changing zone properties

You can change a zone name, a description, the assets assigned to a zone, and the zone set or sets to which a zone belongs.

To change zone properties, perform the following steps:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zones** area, click **Change**.
5. In the Zone Configuration window enter the name of the zone and, optionally, a description. The zone name must be unique in the fabric. Click **Next**.
6. In the next zone set Configuration window, you can change the assets to be included in the zone:
  - In the **Available Assets** column, select any zone assets to be added and click  to move them to the **Selected Assets** column.
  - In the **Selected Assets** column, select any zone assets to be removed from the zone and click  to move them to the **Available Assets** column.
  - Click **Next**.
7. In the next zone set Configuration window, you can change the zone sets to which the zone is assigned:
  - In the **Available Zone Sets** column, select any zone sets to which the zone is to be assigned and click  to move them to the **Selected Zone Sets** column.





- In the **Selected Zone Sets** column, select any zone sets to which the zone is not to be assigned and click  to move them to the **Available Zone Sets** column.

8. Click **Finish**.

## Duplicating zones

You can use the Zones dialog to duplicate a zone.

To duplicate zones, perform the following steps:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zones** area, click **Duplicate**.
5. In the Zone Configuration window enter the name of the zone and, optionally, a description. The zone name must be unique in the fabric. Click **Next**.
6. In the next Zone Configuration window, you can change the assets to be included in the zone:
  - In the **Available Assets** column, select any zones to be added and click  to move them to the **Selected Assets** column. To remove selected zone assets, select the assets from the **Selected Assets** column and click  to move them to the **Available Assets** column.
  - Click **Finish**.
7. In the next Zone Configuration window, you can change the zone sets to which the zone is assigned:
  - In the **Available Zone Sets** column, select any zone sets to which the zone is to be assigned and click  to move them to the **Selected Zone Sets** column.
  - In the **Selected Zone Sets** column, select any zone sets to which the zone is not to be assigned and click  to move them to the **Available Zone Sets** column.
8. Click **Finish**.

## Removing zones

You can use the Zones dialog to remove one or more zones from the list of available zones.

To delete zones, perform the following steps:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zones** area, select one or more zones and click **Remove**. Hold down the Control or Shift key to select multiple zones.
5. In the Zone Deletion window, click **Yes** to delete the selected zones or **No** to cancel the deletion.

## Working with zone sets

This topic describes how to create, change, duplicate, activate or deactivate, and delete zone sets.

## Creating zone sets

This topic provides information about creating new zone sets.

To create a zone set, follow this procedure:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zone Sets** area, click **Add**.
5. Type a name and description for the zone set, and click **Next**. The zone set name must be unique within the fabric.
6. In the zone set Configuration window, do one of the following:
  - Select one or more zones from the **Available Zones** column, move them to the **Selected Zones** column, and click **Finish**.
  - If there are no zones yet defined or if you are not adding any of the zones available at this time, click **Finish**. A zone set Configuration Error window opens stating: "This zone set does not contain any zones. This could result in an error when the zone configuration is applied later. Do you want to continue?" Click **Yes**.


## Changing zone set properties


You can change the zone set definition, which includes the name, description, and the zones that belong to the zone set. When you change and save a zone set definition, the properties are checked by the switches when that definition is stored in the fabric. This can help identify possible problems with that definition before the definition is activated, giving you a chance to correct any problems found before they affect the operation of the fabric.

**Note:** Changing and saving the *definition* for the active zone set does not change the active zone set itself. You must perform an additional step to apply the new definition to the active zone set.

**Note:** After you click **Change** to begin your zoning changes and before you click **Update and Activate** or **Update Only** to conclude your zoning changes, it is possible that other users might have made zoning changes to the same fabric. This may have happened through a switch element manager or switch CLI, as some switch vendors and models lack zone control session locking across the fabric. This could allow overwrite conditions of the changes to the active zoneset for the fabric and to the set of inactive zonesets.

To change the zone set properties, perform the following steps:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zone Sets** area, click the zone set to be changed and click **Change**.
5. In the zone set Configuration window, type a name and description for the zone set, and click **Next**. The zone set name must be unique within the fabric.
6. In the next zone set Configuration window, you can change the zones to be included in the zone set:
  - In the **Available Zones** column, select any zones to be added and click  to move them to the **Selected Zones** column.

- In the **Selected Zones** column, select any zones to be removed from the zone set and click  to move them to the **Available Zones** column.

7. Click **Finish**.

### Activating and deactivating a zone set

Only one zone set can be active at a time. If you deactivate the active zone set, no zone sets are active until you activate another zone set.

To activate or deactivate a zone set, follow these steps:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zone Sets** area, click **Activate** or **Deactivate**.
5. In the **Zone Configuration** area, click one of the following actions:
  - **Update and Activate** to commit the changes you have made to the zone set configuration and immediately activate the zone set.
  - **Update Only** to commit the changes you have made to the zone set configuration without making them active.

### Selecting the inactive zone sets data source

Use the procedure below to change the data source for the inactive zone sets of a selected fabric.

**Note:** This procedure only applies to mixed Brocade and McDATA fabrics that are managed through Proxy SMIS Agent CIMOMs. This procedure does not apply to mixed Brocade and McDATA fabrics that are managed through DCFM Integrated Agent CIMOMs.

To change the data source for an inactive zone set, follow these steps:



1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Select Inactive Zone Sets Data Source**. Use the **Select Inactive Zone Sets Data Source** dialog box to select a data source to be used for gathering and setting the inactive zone sets for the selected fabric. This dialog box has the following fields:
  - **Service URL** displays the CIMOM service URL that can be selected and used for zone control of inactive zone sets for the selected fabric.
  - **Display Name** lists the user-defined display name for the CIMOM.
  - **Connection Status** lists the connection status to the CIMOM.
3. Select a data source and click **OK** to set the data source to be used for probing and controlling inactive zone sets for the fabric. After you select a different data source, a zone probe job is started in the background. A "A probe job for fabric <fabric ID> has been submitted" message box appears. Click **OK** to close the message box.

### Duplicating zone sets

This section provides information about duplicating zone sets.

To create a duplicate zone set, perform the following tasks.

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.

4. In the **Zone Sets** area, click the zone set to be duplicated and click **Duplicate**.
5. In the zone set Configuration window, type a name and description for the zone set, and click **Next**. The zone set name must be unique within the fabric.
6. In the zone set Configuration window, you can change the zones to be included in the new zone set:
  - In the **Available Zones** column, select any zones to be added and click  to move them to the **Selected Zones** column.
  - In the **Selected Zones** column, select any zones to be removed from the zone set and click  to move them to the **Available Zones** column.
7. Click **Finish**.

## Working with zone aliases

This topic describes how to add, change, duplicate, and remove zone aliases. Using zone aliases is a convenient way to work with a named collection of zone members that you want to manage together.

### Creating zone aliases

Creating a zone alias is a convenient way to assign a name to a collection of zone members that you want to manage together.

To create a zone alias:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zones Aliases** area, click **Add**.
5. Enter a unique name for the zone alias that will have zone members. Use the following guidelines for naming a zone alias:
  - a. The maximum length for the zone alias name is 54 characters.
  - b. Allowed characters are: A-Z and a-z; \_ (underscores)
  - c. The name must not begin with a number.
  - d. Spaces are not allowed.
6. Enter an optional description of the zone alias. The description can be any combination of alphanumeric characters.
7. Click **Next**.
8. In the next Zone Alias Configuration window, select the assets (switches, devices, aliases) that you want to include in the zone alias, and move them from the **Available assets** column to the **Selected assets** column, and click **Next**.
9. In the final Zone Alias Configuration window, select one or more zones, move them from the **Available Zones** column to the **Selected Zones** column, and click **Finish**.





**Note:** Not all devices support zone aliases. Please refer to the device manufacturer documentation for further information on zone aliases.

### Changing zone alias properties

You can change a zone alias name, description, the assets assigned to it, and the zone members that are specified in the zone alias.

To change zone alias properties, perform the following steps:







1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zone Aliases** area, select the zone alias you want to change and click **Change**.
5. In the Zone Alias Configuration window change the name of the zone alias and, optionally, enter or change the description. The zone alias name must be unique in the fabric. Use the following naming guidelines:
  - a. The maximum length for the zone alias name is 54 characters.
  - b. Allowed characters are: A-Z and a-z; \_ (underscores)
  - c. The name must not begin with a number.
  - d. Spaces are not allowed.
6. Click **Next**.
7. In the next Zone Alias Configuration window, you can change the assets to be included in the zone alias:
  - In the **Available Assets** column, select one or more assets to be added and click the  to move them to the **Selected Assets** column.
  - In the **Selected Assets** column, select one or more assets to be removed from the zone alias and click the  to move them to the **Available Assets** column.
  - Click **Next**.
8. In the final Zone Alias Configuration window, you can change the zone members as follows:
  - In the **Available Zones** column, select one or more zone members and click the  to move them to the **Selected Zones** column.
  - In the **Selected Zones** column, select one or more zone members and click the  to move them to the **Available Zones** column.
9. Click **Finish**.

### Duplicating zone aliases

To create a new zone alias that's based on an existing one, you can duplicate a zone alias and then make changes to it.

To duplicate an existing zone alias and make changes to it, perform the following steps:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zone Aliases** area, select the zone alias you want to duplicate and click **Duplicate**.
5. In the Zone Alias Configuration window change the name of the zone alias and, optionally, enter or change the description. The zone alias name must be unique in the fabric. Use the following naming guidelines:
  - a. The maximum length for the zone alias name is 54 characters.
  - b. Allowed characters are: A-Z and a-z; \_ (underscores)
  - c. The name must not begin with a number.
  - d. Spaces are not allowed.

6. Click **Next**.
7. In the next Zone Alias Configuration window, you can change the assets to be included in the new zone alias:
  - In the **Available Assets** column, select any assets to be added and click the  to move them to the **Selected Assets** column.
  - In the **Selected Assets** column, select any assets to be removed from the zone alias and click the  to move them to the **Available Assets** column.
  - Click **Next**.
8. In the final Zone Alias Configuration window, you can change the zone alias as follows:
  - In the **Available Zones** column, select one or more zones and click the  to move them to the **Selected Zones** column.
  - In the **Selected Zones** column, select one or more zones and click the  to move them to the **Available Zones** column.
9. Click **Finish**.

### Removing zone aliases

The following procedure removes a selected zone alias from all zones in the currently selected fabric, and removes the zone alias from the fabric's zoning configuration definition. To remove a zone alias from an individual zone or a specific set of zones, use the Zone Configuration wizard.

To remove a zone alias, perform the following steps:

1. In the navigation pane, expand the **Fabric Manager** node and click **Fabrics**.
2. Click a fabric in the list and click **Zone Configuration**.
3. In the **Zone Configuration** area, click **Change**.
4. In the **Zone Aliases** area, select the zone alias you want to remove and click **Remove**. Hold down the Control or Shift key to select multiple zone aliases.
5. In the Zone Alias Deletion window, click **Yes** to delete the selected zone alias.

---

## Managing Tape Libraries

IBM Tivoli Storage Productivity Center provides the capability to manage your tape libraries. For this release, tape library management capabilities are limited to support for the IBM 3584 tape library (full support) and the IBM 3494 tape library (partial support).

You must have at least Tape Operator authority in order to access the functions in the Tape Manager node. To view your authority level, open the "Role-to-Group Mappings" window in the **Administrative Services > Configuration** node in the navigation tree.

1. Discover IBM tape libraries attached to your SAN fabrics.
2. View tape library information.
3. Monitor the performance of groups of tape libraries by probing for data and viewing alerts. .
4. Create and view asset and capacity reports using data collected from your tape libraries.

Asset reports can include such information as the library identifier and the status of the library. Capacity reports include such information as the number of drives in a library, the number of tape slots, and the number of occupied tape slots.

5. Use the topology viewer to view information about tape libraries, as described in “Working with tape libraries in the topology viewer” on page 554.

## Discovering tape libraries


Discovering tape libraries uses the same general process for discovering CIMOMs as discovering storage subsystems does.

- Tape Manager automatically discovers available devices by contacting the CIM agent that is running on those devices.
- You can add Tape CIMOMs by automatic SLP or manually.
- You can create discovery jobs to discover other Tape Library CIM agents, as described in “Discover storage subsystem, tape library, and fabric information with CIM agents” on page 330.
- You can set up your discoveries to run on a regular schedule, for example weekly.

However, you view your tape libraries from a different location in the navigation tree.

- To confirm the tape libraries that were discovered, complete the following steps:
  1. In the navigation tree, expand **Tape Manager**.
  2. Click **Tape Libraries**. The tape libraries that have been discovered will be listed in the right pane.

IBM Tivoli Storage Productivity Center can detect a tape library that is connected to a fabric through queries to the switches in the fabric and through CIMOM queries. However, only basic information is available at this point, and the information from the two discoveries is not correlated. As a consequence, the tape library is displayed in the topology L0:Storage view as two separate icons:

- After the in-band discovery, as a tape library of unknown status (a blue diamond ) identified only by its world-wide name (WWN)
- After the CIMOM discovery, a tape library is identified as, for example, a 3494 or a TS3500

In order to correlate the information collected through the in-band and CIMOM discoveries, you must run a tape library probe. Then the tape library will be displayed by only one icon with its type identified.

## Viewing tape library information

View collected tape library data to stay aware of discovered tape libraries and usage information.

- Click **Tape Manager**. Left-click **Tape Libraries**. In the right pane, the Tape Libraries page is displayed, which lists the discovered tape libraries.  
Select a library from the table and then click one of the buttons to see further details about that library. For example, select a library and click **Drives** to open a page that lists all the drives for that library.

You can also right-click a library in the **Tape Libraries** table and select **View Details**.

- Use the **View Details** feature in the topology viewer, as described in “Working with tape libraries in the topology viewer” on page 554.

## Monitoring tape libraries

Monitor the performance of tape libraries to be aware of events that occur, and to respond to alerts.

1. Create one or more tape library groups, as described in “Creating tape library groups” on page 371.
2. Create a probe job to collect data about your tape libraries and library groups, as described in “Creating tape library probes” on page 338.
3. View alerts that are generated by your tape libraries, as described in “Viewing alert logs” on page 401.

---

## Managing DS8000 storage systems using element managers

Use the Configuration Utility in IBM Tivoli Storage Productivity Center to add and manage element managers for IBM System Storage DS8000. *Element managers* are applications that are used to configure and manage a storage device. You can add and manage connections to the element manager that is provided with DS8000, DS8000 Storage Manager, or you can specify other element managers for a device.

The **Element Manager** tab in the Configuration Utility is used to add and manage DS8000 element managers. The **Element Manager** tab contains multiple sections. The **DS8000 Element Manager** section shows the DS8000 element managers that have been added to Tivoli Storage Productivity Center.

To expand or collapse the **DS8000 Element Manager** section, click the + or - icon to the left of the section heading.

### The DS8000 Element Manager section

The **DS8000 Element Manager** section contains a table listing the element managers that are visible to Tivoli Storage Productivity Center and a **Select Action** list. This list enables you to add an element manager to the section and to manage existing element managers. The menu items that are shown in the **Select Action** list depend on whether you have selected an element manager in the table. This table contains the following fields:

#### Element Manager

Shows the IP address of the element manager or the display name if a display name was defined.

#### Connection Status

Shows the status of the element manager. The values that appear in this field are:

- **Online** (The element manager is accessible.)
- **Offline** (The element manager is not accessible)

#### Device

The name for the device as defined in Tivoli Storage Productivity Center.

#### UID/PWD

Shows whether a user ID and password have been defined for authentication to the element manager. The values that appear in this field are:

- **Yes** (A user ID and password have been defined.)
- **No** (A user ID and password have not been defined.)

**LIC Enabled**

Shows whether a launch descriptor file has been registered for the element manager. A launch descriptor file is an XML file that describes the menus and the detailed uniform resource locators (URLs) for applications. A launch descriptor file is required to enable the launch in context feature for the device. Tivoli Storage Productivity Center provides a launch descriptor file for the element manager that is provided with DS8000.

The values that appear in this field are:

- **Yes** (Launch descriptor file has been registered.)
- **No** (Launch descriptor file has not been registered.)

**Managed**

Shows whether the element manager is associated with a DS8000 that has subsystem interface connection for the device. The values that appear in this field are:

- **Yes** (The element manager is associated with a DS8000 that has subsystem interface connection for the device.)
- **No** (The element manager is not associated with a DS8000 that has subsystem interface connection for the device.)

## Configuring the DS8000 element manager

DS8000 Storage Manager is the element manager that is provided with IBM System Storage DS8000. This section describes how to add and manage DS8000 Storage Manager in IBM Tivoli Storage Productivity Center.

You can perform the following tasks for DS8000 element managers:

- Add DS8000 element managers that you want to manage.
- Start a DS8000 element manager (for example, DS8000 Storage Manager) to configure and administer the storage subsystems associated with the element manager. On Windows, the DS8000 Storage Manager is started in a browser window that is embedded in Tivoli Storage Productivity Center. On AIX, Linux, and UNIX, DS8000 Storage Manager is started in an external Web browser window.
- Modify or remove element managers.
- Test the connection between Tivoli Storage Productivity Center and element managers.

### DS8000 credentials

The configuration of DS8000 credentials depends on whether you are using the single sign-on feature in your environment. *Single sign-on* is an authentication process that enables you to start an application from another application using a single user ID and password. For more information about single sign-on, refer to “Single sign-on” on page 58.

To use single sign-on with DS8000, the DS8000 system must be release 4.2 or later and must be configured for Lightweight Directory Access Protocol (LDAP). Tivoli Storage Productivity Center passes the DS8000 user credentials as a Lightweight Third Party (LTPA) token.

If you are not using single sign-on, the DS8000 element manager user name and password are stored for each Tivoli Storage Productivity Center user. A set of user credentials must be stored for each Tivoli Storage Productivity Center user who

wants to access an DS8000 element manager. The **Element Manager** tab lists all element managers known to Tivoli Storage Productivity Center, regardless of the presence of valid user credentials. Element manager credentials might be out of date or absent for that Tivoli Storage Productivity Center user.

If the element manager has been added to Tivoli Storage Productivity Center using different credentials from the ones you are using, or if you attempt to open the DS8000 element manager for the first time after the element manager was added to Tivoli Storage Productivity Center by DS CIM Agent discovery, the following message is shown:

The element manager GUI cannot be accessed because the user name and password are missing. Update the required fields in the next dialog. The GUI will be displayed after the fields have been correctly completed and the dialog closed.

If you created an element manager with one set of credentials and are currently logged on to Tivoli Storage Productivity Center as another user with different credentials, you can open the DS8000 element manager in one of two ways:

- **Add credentials.** Close the warning message. The **Modify Element Manager** window opens. Add the current credentials and click **Save**. You can then open the DS8000 element manager for that element manager.

**Important:** Use this method with caution because anyone who knows the Tivoli Storage Productivity Center credentials can manage the DS8000 device, including performing destructive operations.

- **Log on as the user who added the element manager.** Close the warning message and the **Modify Element Manager** window. Log off of Tivoli Storage Productivity Center. Log on with the user ID that was used to create the element manager. From the Configuration Utility **Element Manager** tab, start the DS8000 element manager.

## **Adding DS8000 element managers using the Configuration Utility**

Use the Configuration Utility when you want to add one element manager at a time.

To add an element manager in Tivoli Storage Productivity Center, follow these steps:

1. Click **Element Management** in the toolbar, or expand the IBM Tivoli Storage Productivity Center node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab. Any existing element managers are automatically added to the **Element Manager** section when the tab opens.
2. In the **DS8000 Element Manager** section, select **Add Element Manager** from the **Select Action** list. The add Element Manager dialog box displays.
3. Define the information that is required to connect to the element manager using the following fields. Required fields are specified in bold in the dialog box. In the following field descriptions, required fields are marked by an asterisk.

**\*Host** Enter the Domain Name System (DNS) name or Internet Protocol (IP) address of the host computer where the element manager is located. You can enter an IPv4 or IPv6 name or address depending on what address types are configured on the element manager and on the system where Tivoli Storage Productivity Center is installed.

- Enter an IPv4 address in the following format: ###.###.###.###. For example: 192.168.1.65.
- Enter an IPv6 address. These addresses are typically written as eight groups of four hexadecimal digits, but might vary from site to site. For example: 2001:db8:0:0:1812:5b9e:3391:3156. You can use IPv6 addresses if the computer where Tivoli Storage Productivity Center is dual-stack enabled (IPv4 and IPv6).

**Restriction:** If the computer on which the element manager is located is IPv6-only, the Tivoli Storage Productivity Center servers must be installed on an IPv6-only or dual stack (IPv4 and IPv6 addresses are configured) computer.

**\*Port** Enter the port number on which the element manager listens for requests. The default port number that is shown depends on the value that is selected in the **Protocol** field. You can modify the port number if required for your environment.

**Username**

Enter the user name that Tivoli Storage Productivity Center will use to log on to the element manager.

**Password**

Enter the password associated with the user name that is defined in the **Username** field.

If the DS8000 element manager you want to add is a new installation and the enterprise storage server network interface (ESSNI) user password has not been changed from the initial value of **admin**, you are prompted to change the password. Be prepared to enter a password and record the new password in a safe place.

**\*Protocol**

Select the protocol (**http** or **https**) for the element manager.

**Display Name**

Enter a display name for the element manager. The display name is shown in the **Element Manager** field in the DS8000 element manager table. If this field is blank, the IP address of the computer on which the element manager is located is shown in the **Element Manager** field of the table. The display name is used in Tivoli Storage Productivity Center and does not affect the element manager itself.

**Save**

Click this button to add the DS8000 element manager. Tivoli Storage Productivity Center tests the connection to the element manager. If the connection is successful, the element manager is displayed in the DS8000 element manager table. If the connection is unsuccessful, an error message displays indicating that the element manager at the specified IP address could not be added.

**Cancel**

Click this button to close the dialog box without adding the DS8000 element manager.

**Starting DS8000 element managers**

Use the Configuration Utility to start IBM System Storage DS8000 element managers from IBM Tivoli Storage Productivity Center. You can also start DS8000 element managers from other points in Tivoli Storage Productivity Center.

To start DS8000 element managers using the Configuration Utility follow these steps:

1. Expand the **IBM Tivoli Storage Productivity Center** node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab. **Note:** You can also access the Element Manager tab by selecting **Element Management** in the tool bar.
2. In the **DS8000 Element Manager** section, click the element manager that you want to start from the DS8000 element manager table.
3. Use one of the following methods to start the element manager:
  - To start the default element manager for DS8000, select **Launch Default Element Manager** from the **Select Action** list.
  - To start another element manager for DS8000, right click the element manager that you want to start and click **Launch Element Manager**. Select the element manager that you want to start from the menu.

If the connection information for an element manager is valid, the element manager is started. If unsuccessful, an error message is displayed. The error messages that you receive describes one of the following problems:

- that the element manager at the specified IP address could not be started
- that the required connection information for the element manager has not been provided
- that the element manager is not installed and configured correctly (for local element manager installations such as XIV Storage Manager)

If you are prompted to enter additional information such as an ID and password for the element manager, complete the following steps:

1. Click **OK** in the message dialog box. The Modify Element Manager dialog box is displayed.
2. Enter the required information for the element manager and click **Save**.

## Modifying DS8000 element managers

Use the Configuration Utility to modify information for IBM System Storage DS8000 element managers in IBM Tivoli Storage Productivity Center.

To modify an element manager in Tivoli Storage Productivity Center, follow these steps:

1. Click **Element Management** in the toolbar, or expand the IBM Tivoli Storage Productivity Center node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **DS8000 Element Manager** section, select the element manager that you want to modify from the DS8000 element manager table.
3. Select **Modify Element Manager** from the **Select Action** list. The Modify Element Manager dialog box is displayed.
4. Modify any of the following fields:

### Username

Enter the user name that Tivoli Storage Productivity Center will use to log on to the element manager.

### Password

Enter the password associated with the user name that is defined in the **Username** field.



If the DS8000 element manager you want to add is a new installation and the enterprise storage server network interface (ESSNI) user password has not been changed from the initial value of **admin**, you are prompted to change the password. Be prepared to enter a password and record the new password in a safe place.

**Password Confirm**

Enter the same password that you entered in the **Password** field.

**Display Name**

Enter a display name for the element manager. The display name is shown in the **Element Manager** field in the DS8000 element manager table. If this field is blank, the IP address of the computer on which the element manager is located is shown in the **Element Manager** field of the table. The display name is used in Tivoli Storage Productivity Center and does not affect the element manager itself.

**Save** Click this button to modify the element manager information. Tivoli Storage Productivity Center tests the connection to the element manager. If the connection is successful, the element manager information is saved. If the connection is unsuccessful, an error message appears indicating that the element manager at the specified IP address could not be added.

**Cancel**

Click this button to close the dialog box without modifying the element manager information

## Removing DS8000 element managers

Use the Configuration Utility to remove IBM System Storage DS8000 element managers from IBM Tivoli Storage Productivity Center. This action does not affect DS8000 Storage Manager on the Hardware Management Console (HMC).

To remove an element manager from Tivoli Storage Productivity Center, complete the following steps:

1. Click **Element Management** in the toolbar, or expand the IBM Tivoli Storage Productivity Center node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **DS8000 Element Manager** section, select the element manager that you want to remove from the DS8000 element manager table.
3. Select **Remove Element Manager** from the **Select Action** list. A confirmation dialog box is displayed.
4. Click **OK**. The element manager is deleted from the DS8000 element manager table. When you remove an element manager, it is removed only from the list of element managers that are visible to Tivoli Storage Productivity Center.

To exit the window without removing the element manager, click **Cancel**.

You can remove only one element manager from the DS8000 element manager table at a time.

## Testing the connection to DS8000 element managers

Use the Configuration Utility to test the connection to DS8000 element managers from IBM Tivoli Storage Productivity Center.

To test the connection to an element manager from Tivoli Storage Productivity Center, follow these steps:

1. Click **Element Management** in the toolbar, or expand the IBM Tivoli Storage Productivity Center node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **DS8000 Element Manager** section, select the element manager for which you want to test the connection from the DS8000 element manager table.
3. Select **Test Element Manager Connection** from the **Select Action** list. IBM Tivoli Storage Productivity Center tests the status of the element manager connection. A message dialog box that states whether the test passed or failed is displayed.

## Configuring other element managers for DS8000

DS8000 Storage Manager is the element manager that is provided with IBM System Storage DS8000. However, you can add other element managers for DS8000 as described in this section.

### Adding user-defined element managers for DS8000

A *user-defined element manager* for IBM System Storage DS8000 is an element manager other than DS8000 Storage Manager. Use the Configuration Utility to add user-defined element managers for IBM System Storage DS8000 in IBM Tivoli Storage Productivity Center.

User-defined element managers are not displayed in the **DS8000 Element Manager** table. These element managers are displayed in the various launch menus that are provided in IBM Tivoli Storage Productivity Center.

To add user-defined element managers in Tivoli Storage Productivity Center, follow these steps:

1. In the navigation tree, expand **IBM Tivoli Storage Productivity Center**. Left-click **Configuration Utility**. In the right pane, click the **Element Manager** tab.
2. In the **DS8000 Element Manager** section, select the row in the DS8000 element manager table that contains the device to which you want to add a user-defined element manager.
3. Select **Add User Defined Element Manager** from the **Select Action** list. The Add User Defined Element Manager dialog box is displayed.
4. Enter a Uniform Resource Locator (URL) or command to start the element manager, and click **Save**.

To exit without adding a element manager, click **Cancel**.

### Removing user-defined element managers for DS8000

Use the Configuration Utility to remove user-defined element managers for IBM System Storage DS8000 in IBM Tivoli Storage Productivity Center.

To remove a user-defined element manager from Tivoli Storage Productivity Center, follow these steps:

1. Click **Element Management** in the toolbar, or expand the IBM Tivoli Storage Productivity Center node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **DS8000 Element Manager** section, select the row in the DS8000 element manager table that contains the device from which you want to remove a user-defined element manager.
3. Select **Remove Element Manager** from the **Select Action** list. The Remove Element Manager dialog box opens.

4. Select the element manager that you want to remove from the list and click **OK**. The element manager is removed. When you remove an element manager, it is removed only from the list of element managers that are visible to Tivoli Storage Productivity Center.

To exit the window without removing the element manager, click **Cancel**.

---

## Managing other storage systems using element managers

Use the Configuration Utility in IBM Tivoli Storage Productivity Center to add and manage element managers for storage devices other than IBM System Storage DS8000. *Element managers* are applications that are used to configure and manage a storage device. You manage connections to the element managers that are provided with the storage device or you can specify other element managers for a device.

The **Element Manager** tab in the Configuration Utility is used to add and manage element managers. The **Element Manager** tab contains multiple sections. The **Other Device Element Manager** section shows element managers that have been added using a discovery job or existing element managers that were added to Tivoli Storage Productivity Center prior to release 4.1.

If you are gathering data from IBM XIV Storage System, the XIV system element manager is displayed in the **Other Device Element Manager** section. However, the tasks available for the other element managers in this section are not available for the XIV system element manager. Use the Storage Subsystem page to work with the XIV system element manager (see “Managing storage subsystems” on page 565).

To expand or collapse the **Other Device Element Manager** section, click the + or - icon to the left of the section heading.

### Other Device Element Manager section

The **Other Device Element Manager** section contains a table listing the element managers or external tools that are visible to Tivoli Storage Productivity Center and a **Select Action** list. This list enables you to manage existing element managers. The menu items that are shown in the **Select Action** list depend on whether you have selected an element manager in the table. This table contains the following fields:

#### Element Manager

Shows the IP address of the element manager or the display name if a display name was defined.

#### Connection Status

Shows the status of the element manager. The values that appear in this field are:

- **Online** (The element manager is accessible.)
- **Offline** (The element manager is not accessible)

#### Device

The name for the device as defined in Tivoli Storage Productivity Center.

#### UID/PWD

Shows whether a user ID and password have been defined for authentication to the element manager. The values that appear in this field are:

- **Yes** (A user ID and password have been defined.)

- **No** (A user ID and password have not been defined.)

#### **LIC Enabled**

Shows whether a launch descriptor file has been registered for the element manager. A launch descriptor file is an XML file that describes the menus and the detailed uniform resource locators (URLs) for applications. A launch descriptor file is required to enable the launch in context feature for the device.

The values that appear in this field are:

- **Yes** (Launch descriptor file has been registered.)
- **No** (Launch descriptor file has not been registered.)

## **Adding element managers using CIMOM discovery**

Use the CIM object manager (CIMOM) discovery process in IBM Tivoli Storage Productivity Center to add element managers that are managing the same storage subsystems as CIMOMs in your environment. The element managers discovered through the CIMOM discovery process are automatically added to the **Element Manager** tab in the Configuration Utility.

Use this method when you want to add multiple element managers to Tivoli Storage Productivity Center simultaneously. For example, use these steps when:

- You want to run a discovery against multiple CIMOMs and each of those CIMOMs has a registered storage system.
- You have more than one registered storage system per CIMOM.

This method does not apply to storage systems that use native interfaces: IBM System Storage DS8000 Storage Manager, IBM XIV Storage System, IBM System Storage SAN Volume Controller, IBM Storwize V7000 Unified, and IBM Storwize V7000.

This procedure does not physically add the CIMOM to your environment, but provides Tivoli Storage Productivity Center with the information it needs to communicate with that CIMOM.

To add an element manager using CIMOM discovery in Tivoli Storage Productivity Center, complete the following steps:

1. Expand **Administrative Services > Data Sources > CIMOM Agents** in the navigation tree.
2. In the right pane, click **Add CIMOM**. The Add CIMOM dialog box is displayed.
3. Define the information that is required to add the CIMOM using the following fields. Required fields are specified in bold in the dialog box. In the following field descriptions, required fields are marked by an asterisk.

**\*Host** Enter the IP address of the computer on which the CIMOM is installed. You can enter an IPv4 or IPv6 address depending on what is supported on the computer where the CIMOM is located:

- Enter an IPv4 address in the following format: **###.###.###.###**. For example: 192.168.1.65
- Enter an IPv6 address. These addresses are typically written as eight groups of four hexadecimal digits, but might vary from site to site. For example: 2001:DB8::/32:1812:5b9e:3391:3156. You can use IPv6 addresses if the computer where Tivoli Storage Productivity Center is installed is IPv6 or dual stack (IPv4 and IPv6) enabled.

**Note:** If the computer on which the CIMOM is located is IPv6-only, the Tivoli Storage Productivity Center servers must be installed on an IPv6-only or dual stack (IPv4 and IPv6 enabled) computer.

**Port** Enter the port through which Tivoli Storage Productivity Center contacts the CIMOM. You can modify the default port number if required for your environment.

**Username**

If required by the CIMOM, enter the user name that Tivoli Storage Productivity Center uses to log on to the CIMOM.

**Password**

Enter the password for the user name that is defined in the **Username** field.

**\*Interoperability Namespace**

Enter the interoperability namespace of the CIMOM. This namespace within the CIMOM allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how Tivoli Storage Productivity Center interacts with the CIMOM when retrieving information.

**Protocol**

Select the protocol (**http** or **https**) for the CIMOM.

**Truststore Location**

Enter the location (path on this computer) of the truststore for the CIMOM. The truststore is the certificate file for certificate-based authentication in the https protocol.

**Truststore Passphrase**

Enter the passphrase of the truststore for the CIMOM.

**Display Name**

Enter the name for the CIMOM as you want it to be shown in Tivoli Storage Productivity Center.

**Description**

Enter a short description for the CIMOM as you want it to be shown in Tivoli Storage Productivity Center.

**Test CIMOM connectivity before adding**

Select this check box to ensure that communication to the CIMOM is working properly before the CIMOM is added. If this box is not checked, the CIMOM is added without verifying connectivity with the given parameters.

**Important:** Failure to check the CIMOM connection might cause a failure later if the CIMOM cannot be scanned or probed.

**Save** Click this button add the CIMOM.

- If you selected **Test CIMOM connectivity before adding**, Tivoli Storage Productivity Center attempts to connect to the CIMOM using the information provided in this dialog box. If the logon information for a CIMOM is incorrect, an error occurs and an associated error message dialog box is displayed.
- If you did not select **Test CIMOM connectivity before adding**, the CIMOM is associated with the element manager but connectivity with Tivoli Storage Productivity Center is not verified.

**Cancel**

Click this button to close the dialog box without adding the CIMOM.

4. Expand **Administrative Services > Discovery** in the navigation tree.
5. Right-click **CIMOM** and select **Run Now**. The CIMOM discovery process adds the element managers that are managing the same storage system as the CIMOM. The element managers are displayed on the appropriate element manager section on the Element Manager tab of the Configuration Utility. Element managers that are added through CIMOM discovery might exist on a Hardware Management Console (HMC) other than the HMC where the CIMOM is located.

## Modifying other element managers

Use the Configuration Utility to modify information for element managers in IBM Tivoli Storage Productivity Center.

To modify an element manager other than an IBM System Storage DS8000 element manager in Tivoli Storage Productivity Center, follow these steps:

1. Click **Element Management** in the toolbar, or expand the IBM Tivoli Storage Productivity Center node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **Other Device Element Manager** section, select the element manager that you want to modify from the element manager table.
3. Select **Modify Element Manager** from the **Select Action** list. The Modify Element Manager dialog box is displayed.
4. Modify any of the following fields:

**Username**

Enter the user name that Tivoli Storage Productivity Center will use to log on to the element manager.

**Password**

Enter the password associated with the user name that is defined in the **Username** field.

**Password Confirm**

Enter the same password that you entered in the **Password** field.

**Display Name**

Enter a display name for the element manager. The display name is shown in the **Element Manager** field in the element manager table. If this field is blank, the IP address of the computer on which the element manager is located is shown in the **Element Manager** field of the table. The display name is used in Tivoli Storage Productivity Center and does not affect the element manager itself.

**Save**

Click this button to modify the element manager information. Tivoli Storage Productivity Center tests the connection to the element manager. If the connection is successful, the element manager information is saved. If the connection is unsuccessful, an error message appears indicating that the element manager at the specified IP address could not be added.

**Cancel**

Click this button to close the dialog box without modifying the element manager information

## Starting other element managers

Use the Configuration Utility to start element managers from IBM Tivoli Storage Productivity Center. You can also start element managers from other points in Tivoli Storage Productivity Center.

To start element managers other than IBM System Storage DS8000 element managers from the Configuration Utility follow these steps:

1. Click **Element Management** in the toolbar, or expand the IBM Tivoli Storage Productivity Center node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **Other Device Element Manager** section, click the element manager that you want to start from the element manager table.
3. Use one of the following methods to start the element manager:
  - To start the default element manager for a storage system, select **Launch Default Element Manager** from the **Select Action** list.
  - To start another element manager for a storage system, right click the element manager that you want to start and click **Launch Element Manager**. Select the element manager that you want to start from the menu.

If the connection information for an element manager is valid, the element manager is started. If unsuccessful, an error message is displayed. The error message that you receive describes one of the following problems:

- that the element manager at the specified IP address could not be started
- that the required connection information for the element manager has not been provided
- that the element manager is not installed and configured correctly (for local element manager installations such as XIV Storage Manager)

If you are prompted to enter additional information such as an ID and password for the element manager, complete the following steps:

1. Click **OK** in the message dialog box. The Modify Element Manager dialog box is displayed.
2. Enter the required information for the element manager and click **Save**.

## Testing the connection to other element managers

Use the Configuration Utility to test the connection to element managers from IBM Tivoli Storage Productivity Center.

To test the connection to an element manager other than an IBM System Storage DS8000 element manager from Tivoli Storage Productivity Center, follow these steps:

1. Click **Element Management** in the toolbar, or expand the IBM Tivoli Storage Productivity Center node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **Other Device Element Manager** section, select the element manager for which you want to test the connection from the element manager table.
3. Select **Test Element Manager Connection** from the **Select Action** list. IBM Tivoli Storage Productivity Center tests the status of the element manager connection. A message dialog box that states whether the test passed or failed is displayed.

## Adding user-defined element managers for other storage systems

A *user-defined element manager* is an element manager other than the manager that is provided with the storage system. Use the Configuration Utility to add user-defined element managers in IBM Tivoli Storage Productivity Center.

User-defined element managers are not displayed in the **Other Device Element Manager** table. These element managers are displayed in the various launch menus that are provided in IBM Tivoli Storage Productivity Center.

To add user-defined element managers for a storage device other than DS8000 in Tivoli Storage Productivity Center, follow these steps:

1. Click **Element Management** in the toolbar, or expand the IBM Tivoli Storage Productivity Center node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **Other Device Element Manager** section, select the row in the element manager table that contains the storage system to which you want to add a user-defined element manager.
3. Select **Add User Defined Element Manager** from the **Select Action** list. The Add User Defined Element Manager dialog box is displayed.
4. Enter a Uniform Resource Locator (URL) or command to start the element manager, and click **Save**.

To exit without adding a element manager, click **Cancel**.

## Removing user-defined element managers for other storage systems

Use the Configuration Utility to remove user-defined element managers for storage systems in IBM Tivoli Storage Productivity Center.

To remove a user-defined element managers for a storage device other than DS8000 in Tivoli Storage Productivity Center, follow these steps:

1. Click **Element Management** in the toolbar, or expand the IBM Tivoli Storage Productivity Center node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **Other Device Element Manager** section, select the row in the element manager table that contains the storage system from which you want to remove a user-defined element manager.
3. Select **Remove Element Manager** from the **Select Action** list. The Remove Element Manager dialog box opens.
4. Select the element manager that you want to remove from the list and click **OK**. The element manager is removed. When you remove an element manager, it is removed only from the list of element managers that are visible to Tivoli Storage Productivity Center.

To exit the window without removing the element manager, click **Cancel**.

---

## Launch and launch in context

The launch and launch in context features are used to start IBM Tivoli Storage Productivity Center and to start other applications from Tivoli Storage Productivity Center. The terms inbound and outbound are used to differentiate starting Tivoli Storage Productivity Center (inbound) and starting another application from Tivoli Storage Productivity Center (outbound).



Inbound launch in context refers to starting Tivoli Storage Productivity Center at a specific point in the graphical user interface (GUI). You can use inbound launch in context in either of the following ways:

- Configure an application such as Tivoli Enterprise Portal to start Tivoli Storage Productivity Center at functionally related locations in the Tivoli Storage Productivity Center GUI.
- Start Tivoli Storage Productivity Center from Java Web Start, the command line, or the Windows Start menu using parameters to specify the starting location in the Tivoli Storage Productivity Center GUI.

Outbound launch in context refers to starting other applications, such as administrative applications for storage systems and switches, from functionally related points in the Tivoli Storage Productivity Center GUI.

Launch refers to starting an application at the starting point for that application. Launch primarily refers to starting applications from the Tivoli Storage Productivity Center GUI. Starting applications from Tivoli Storage Productivity Center is referred to as an outbound launch.

**Related concepts:**

“Inbound launch in context feature” on page 61

With the inbound launch in context feature you can start IBM Tivoli Storage Productivity Center at a specified point in the graphical user interface (GUI).

“**Outbound launch and launch in context features**” on page 67

With the outbound launch and launch in context features, you can start the administrative applications for devices such as storage systems and switches from various points in the graphical user interface (GUI).

## Inbound launch in context feature

With the inbound launch in context feature you can start IBM Tivoli Storage Productivity Center at a specified point in the graphical user interface (GUI).

**Prerequisite:** You must have View and Modify permissions to use the launch in context feature.

You can use the launch in context feature to perform the following tasks in Tivoli Storage Productivity Center:

- Create a volume or VDisk
- Open the storage area network SAN Planner
- Create a capacity report
- Create a wasted space report

You can run the launch in context feature in the following ways:

- **Use Java Web Start** from a Web browser. See “Launch with Java Web Start” on page 61.
- **Use the command line on Windows, AIX, Linux, or UNIX** to enter the command and arguments on the command line. See “Launch from the command line” on page 63.
- **Use the Windows Start menu.** See “Launch from Start menu” on page 64.
- **Use the Tivoli Enterprise Portal GUI** to define, edit, and launch commands. See “Launch from Tivoli Enterprise Portal” on page 65.

The launch in context feature has a number of parameters. See “Launch in context parameters” on page 1151, “Launch in context task parameters” on page 1153, and “Available windows” on page 1154.

## Launch with Java Web Start

Using Java Web Start, you can use the launch in context feature from a Uniform Resource Locator (URL).

The launch in context feature requires Java Web Start 6 on the computer that is running the Web browser. If the correct release of Java Web Start is not present, you receive an error message when you enter the URL. The message states that you must download the Java Runtime Environment, which includes Java Web Start. Follow the link provided in the message to install the required software.

If you are using Internet Protocol Version 6 (IPv6) and are using Java Web Start to start IBM Tivoli Storage Productivity Center on a Windows system, use a system host name such as **tpc332.storage.newyork.xyz.com:9549** in the Web browser address field rather than an IPv6 address such as **[2001:DB8:0000:1234:0000:0000:5678:ABCD]:9549**.

To use the launch in context feature to perform a task or display a window on a remote server from your browser, complete the following steps:

1. Open your browser and type the launch in context command.

URLs cannot include spaces. In the command, a parameter and its value are separated by an equal sign (=). Parameter=value pairs are separated by an ampersand (&). If any of the parameters are spelled incorrectly, an error message indicates the incorrect parameter.

The format of a command to create a volume is:

```
https://device_server:device_port/ITSRM/jsp/jnlp/tpcgui.jsp?consoleClass=com.ibm.usmi.console.lic.tpc.LICConsole&-server=server:port&-function=disk.create_volume&-subsystem_id=ESS-2105-17844-IBM&-name=myvol&-number=2&-size=10&-unit=MB
```

The format of a command to view a window on a remote server is:

```
http://device_server:device_port/ITSRM/jsp/jnlp/tpcgui.jsp?consoleClass=com.ibm.usmi.console.lic.tpc.LICConsole&-server=server:port&-function=window_name
```

**Tip:** The Data server port value is 9549. The port value for the Device server is 9551 for https and 9550 for http. If the wrong the Device server port value is used, the return page cannot be interpreted by the browser. If the wrong Data server port value is used, the following error message is returned.

Unable to connect to the server

2. The first time a window or task is started, you receive a security warning informing you that the downloaded jar files have been digitally signed using a certificate from a trusted company. To prevent this message from being displayed in the future, select **Always trust content from this publisher** and click **Run**. If the Tivoli Storage Productivity Center graphical user interface (GUI) has not been downloaded on the local computer using Java Web Start, or the local computer does not have the latest version of the Tivoli Storage Productivity Center GUI, the GUI is downloaded.
3. If you included the user ID and password in the command, the sign-in dialog is bypassed, otherwise log on to the Tivoli Storage Productivity Center GUI. The Tivoli Storage Productivity Center GUI navigates to the tree node corresponding to the window name that you specified. You are connected using

the user name and password that you entered. Depending on the value you specified for the `-function` parameter, the following actions might occur:

- The specified window opens
- For `disk.create_volume`, depending on the parameters, the Create Volume wizard or Create VDisk wizard opens.
- For `disk.san_planner`, the SAN Planner panel opens with the given SAN plan ID.
- For `data.filesystem_capacity_by_computer`, the Capacity Report opens.
- For `TPC.reports_data_wasted_space`, the Wasted Space Report opens.

#### Restrictions:

- When you open a plan in the SAN Planner using launch in context, you must enter the exact name of the plan. The value for the `-plan_id` parameter is case-sensitive and must match the name of the plan when it was created. For example, if you created a plan named **UserPlan1**, you must specify `-plan_id UserPlan1` in the launch in context command.
- When you launch the Filesystem Capacity > By Computer report using launch in context, you must specify the ID for the computer upon which you want to report. For example:

```
http://tpcserver14:9550/ITSRM/jsp/jnlp/tpcgui.jsp?consoleClass=com.ibm.usmi.console.lic.tpc.LICConsole&-server=tpcserver14.storage.xyz.com:9549&-function=data.filesystem_capacity_by_computer&-computer_id=tpcserver14
```

The value you enter for the `computer_id` parameter might not match the ID for that computer that is stored in the Tivoli Storage Productivity Center database repository. You receive a warning message if the ID for the computer that you enter in the URL does not match the ID that is stored in the data repository. To ensure that this report is displayed properly when you run from the launch in context facility, enter an ID for the computer that matches the ID that is displayed for it through the Topology Viewer or appropriate report in the product user interface.

Use the information in the links below to enter parameters.

“Launch in context parameters” on page 1151

You can use the launch in context parameters, definitions, and examples.

“Launch in context task parameters” on page 1153

Use the launch in context feature task parameters to perform tasks.

“Available windows” on page 1154

You can use windows with the inbound launch in context feature to start IBM Tivoli Storage Productivity Center. These windows are called with the `-function window name` argument.

### Launch from the command line

This topic describes how to start windows and tasks from the Windows AIX, Linux, or UNIX command line.

The bat file `tpc.bat` is used for Windows operating systems. The `.sh` file `TPCD.sh` is used for AIX, Linux, or UNIX systems. The parameters follow the file name. Parameters are separated by spaces. Parameters are optional and can appear in any order. See “Launch in context parameters” on page 1151, “Launch in context task parameters” on page 1153, and “Available windows” on page 1154 for the parameters.

Frequently used commands can be saved as unique command files and rerun as needed.

1. In the Command Prompt window, type the command for the action that you want to take, and press Enter.

The action occurs on the remote Data server. The graphical user interface (GUI) window for that action opens on the local computer.

2. Enter parameters if you previously did not enter any parameters. If you do not provide a user name, password, or server name, the IBM Tivoli Storage Productivity Center GUI opens the connection dialog box. If you do not provide these values, the GUI presents an error message box and requests the values. If you specify incorrect parameters, you receive an error message.

The following example launches a task from the Windows command line:

```
tpc.bat -user admin -passwd password1 -server
tivoli122.storage.usca.ibm.com:9549 -function
data.filesystem_capacity_by_computer -computer_id <computerID>
```

The following example uses the launch in context feature to view a window on the remote machine:

```
tpc.bat -user admin -passwd password1 -server
tivoli122.storage.usca.ibm.com:9549 -function disk.storage_subsystem
```

#### Restrictions:

- When you open a plan in the SAN Planner using the command line or a Web browser for launch in context, you must enter the exact name of the plan. The value for **-plan\_id** in the command line is case-sensitive and must match the name of the plan when it was created. For example, if you created a plan named **UserPlan1**, you might enter the following command on the command line:

```
./TPCD.sh -user administrator -passwd passw0rd -server
tpcserver14.storage.xyz.com:9549 -function disk.san_Planner -plan_id
UserPlan1
```

- There is a known limitation when you start the Filesystem Capacity > By Computer report using launch in context. To start this report, you must specify the ID for the computer upon which you want to report. For example:

```
./TPCD.sh -user root -passwd "passw0rd" -server
tpcserver14.storage.xyz.com:9549 -function
data.filesystem_capacity_by_computer -computer_id tpcserver14
```

The value you enter for "computer\_id" might not match the ID for that computer that is stored in the Tivoli Storage Productivity Center database repository. You will receive a warning message if the ID for the computer you enter in the command line does not match the ID stored in the data repository. To ensure that this report is displayed properly when run from the launch in context command line or through Java Web Start, make sure to enter an ID for the computer that matches the ID that is displayed for it through the Topology Viewer or appropriate report in the product's user interface.

## Launch from Start menu

This topic describes how to start windows and actions from the Windows Start menu.

The bat file, `tpc.bat`, is used for Windows operating systems. The parameters follow the bat file name. Parameters are optional and can appear in any order. To use the launch in context feature from the Windows Start menu, complete the following steps:

1. Click **Start > Programs > IBM Tivoli Storage Productivity Center**, right-click **Productivity Center** and select **Properties**. The Productivity Center Properties window opens.
2. Enter the command-line arguments in the **Target** text field and click **Enter**.  
If the Data server is remote, the action occurs on the Data server and the graphical user interface (GUI) window the action opens is on the local computer.
3. Enter the parameters if you did not previously enter the necessary parameters. If you do not provide the user name, password, or server name, the Tivoli Storage Productivity Center GUI opens the connection dialog box. If you do not provide these values, the GUI presents an error message box and requests the values.

The following example starts a task from the Windows Start menu:

```
"C:\Program Files\IBM\TPC\gui\tpc.bat" -user admin -passwd password1
-server tivoli122.storage.usca.ibm.com:9549 -function
data.filesystem_capacity_by_computer -computer_id computerID
```

The following example uses the Windows Start menu to view a window on the remote server:

```
"C:\Program Files\IBM\TPC\gui\tpc.bat" -user admin -passwd password1
-server tivoli122.storage.usca.ibm.com:9549 -function
disk.storage_subsystem
```

#### For more information see:

“Launch in context parameters” on page 1151

You can use the launch in context parameters, definitions, and examples.

“Launch in context task parameters” on page 1153

Use the launch in context feature task parameters to perform tasks.

“Available windows” on page 1154

You can use windows with the inbound launch in context feature to start IBM Tivoli Storage Productivity Center. These windows are called with the `-function window name` argument.

## Launch from Tivoli Enterprise Portal

This topic shows you how to start an existing launch definition from the Tivoli Enterprise Portal graphical user interface (GUI).

You must configure Tivoli Enterprise Portal before you can use the launch in context feature. For information on prerequisites and details on how to configure the launch in context feature with Tivoli Enterprise Portal, see “Configuring Tivoli Enterprise Portal” on page 65. After Tivoli Enterprise Portal has been configured, you can start the application from the configuration window or use **Launch** from the Tivoli Enterprise Portal navigation tree.

If you use the Tivoli Enterprise Portal navigation tree, **Launch** will open the Tivoli Enterprise Portal graphical user interface with your predefined launch in context commands. To use one of these commands, complete the following steps:

1. In the Tivoli Enterprise Portal window, expand **Enterprise > Windows Systems** *your\_server\_name* > **Windows OS**. Right-click **Disk** and select **Launch**. Disk is one of several values you can select. Values include Disk, Network, Enterprise Services, Memory, Printer, Process, Processor, and System.
2. The Create or Edit Launch Definitions window opens showing a list of defined launch definitions in the **Existing Launch Definitions** area.
3. Select the launch definition you want to use from the list and click **Launch**. The window or task that is defined by that launch definition opens.

For information on editing the launch in context command with the Tivoli Enterprise Portal GUI, see:

**Related tasks:**

“Editing launch definitions” on page 67

You can edit an existing launch definition with the Tivoli Enterprise Portal graphical user interface (GUI).

**Configuring Tivoli Enterprise Portal:**

You can configure Tivoli Enterprise Portal so that you can use the launch in context feature.

To configure Tivoli Enterprise Portal complete the following steps:

1. The initial steps depend on whether you are going to open a window or perform a task.
  - a. If you are opening a window, in the Tivoli Enterprise Portal navigation window on the local computer, select where you want the launch definition to originate from. This opens the corresponding view on the target computer. Right-click one of the following locations.
    - The navigator item
    - A row in a table or situation event console view
    - A slice of a pie chart
    - A bar of a bar chart

If the launch definition originates from a table, chart, or situation event console view, open the workspace containing the view. If you think you might edit the argument later, create the launch definition from a table or chart view.
  - b. If you are going to perform a task, create a launch definition for one of the following tasks:
    - Create a volume or VDisk
    - Open the SAN Planner
    - Create a capacity report
    - Create a wasted space report
2. In the Tivoli Enterprise Portal window, expand **Enterprise > Windows Systems** *your\_server\_name* > **Windows OS**. Right-click **Disk** and select **Launch**. Disk is one of several values you can select. Values include Disk, Network, Enterprise Services, Memory, Printer, Process, Processor, and System.
3. The Create or Edit Launch Definitions window opens. Click **Create New**.
4. In the **Name** text field, type a descriptive name for this launch definition. This name is used in the definition list.

5. In the **Target** text box click **Browse** to locate the tpc.bat file. The Select files for directories for Launch Definitions window opens. A typical location for Windows systems is C:\Program Files\IBM\TPC\gui\tpc.bat. Select the correct directory and click **OK**.

When you start a launch definition from the command line, if any part of the path contains spaces, enclose the entire entry in quotation marks.

6. In the **Arguments** text field, enter any parameters that will be added when the application is started. Arguments are entered as *-parameter value*. In the Create or Edit Launch Definitions window, click **Browse** next to **Arguments** to locate the parameters that you want to use. The Select Browser Options window opens. Click **Substitutable Items** and then click **OK**. Select parameters from the **Selected Context** and **Extended Context** lists.

The **help** argument lists the available **function** values. You must type the help parameter by itself into the **Arguments** text field or into the form tpc.bat -help functions on the command line for Windows or TPCD.sh -help functions for UNIX. The **-help functions** parameter provides a window that lists the nodes in the navigation tree that are available to start. If you type **-help** without the **functions** parameter, it results in an error message.

When all parameters have been selected, click **OK** to return to the Create or Edit Launch Definitions window.

7. If the parameters can be changed at launch time (for example, you want to change the system ID), check the **Edit arguments before execution** box.
8. Specify a value for the **Start in** location. Click **Browse** to locate the IBM Tivoli Storage Productivity Center installation and select the folder **gui** where tpc.bat is located.
9. Click **Evaluate** to see how the command resolves when the application starts. Correct any of the parameters that are incorrect and reevaluate the command.
10. When you are satisfied with the launch definition, click **Apply** to save your changes. This keeps the window open so that you can start the current launch definition or create a new one. Click **OK** to save the launch definition and close the window.
11. To start the current definition, select **Launch**. The Tivoli Storage Productivity Center graphical user interface (GUI) with the function parameter that you have provided is opened.

#### **Related reference:**

“Launch in context parameters” on page 1151

You can use the launch in context parameters, definitions, and examples.

“Launch in context task parameters” on page 1153

Use the launch in context feature task parameters to perform tasks.

“Available windows” on page 1154

You can use windows with the inbound launch in context feature to start IBM Tivoli Storage Productivity Center. These windows are called with the *-function window name* argument.

#### **Editing launch definitions:**

You can edit an existing launch definition with the Tivoli Enterprise Portal graphical user interface (GUI).

If the **Edit arguments before execution** box was checked when a launch definition was configured, you can edit the launch definition by changing the attribute values. To edit a launch definition, complete the following steps:

1. In the Tivoli Enterprise Portal window, expand **Enterprise > Windows Systems > your\_server\_name > Windows OS**. Right-click **Disk** and select **Launch**. Disk is one of several nodes you can select. Nodes include Disk, Network, Enterprise Services, Memory, Printer, Process, Processor, and System.
2. The Create or Edit Launch Definitions window opens showing a list of launch definitions in the **Existing Launch Definitions** area.
3. Select the launch definition that you want to edit. The values are displayed.
4. Make the necessary changes in the text fields.
5. Click **Evaluate** to see how the arguments will resolve when the application starts. Correct any incorrect arguments and re-evaluate the command.
6. Click **Apply** to save the changes and keep the window open. Click **OK** to save the launch definition and close the window.  
If you do not want to overwrite the existing launch definition, save the new launch definition with a unique name.

## Outbound launch and launch in context features

With the outbound launch and launch in context features, you can start the administrative applications for devices such as storage systems and switches from various points in the graphical user interface (GUI).

There are two options for starting an application from Tivoli Storage Productivity Center. The option that is available depends on the device that you want to start.

The launch feature starts an application at the starting point for that application.

The launch in context feature starts an application at a functionally related location in that application. Launch in context is available for the following applications:

- Tivoli Storage Productivity Center for Replication
- SAN Volume Controller 6.1 or later
- Storwize V7000
- Storwize V7000 Unified










### Launch menu

When you right-click one of the following entities in Tivoli Storage Productivity Center GUI, you see a launch menu for that entity. The content of the launch menu varies as described in “Launch menu items” on page 69.

- **Alert Log Viewer**. Each row of the alert log is enabled for a launch in context menu. The context is determined by the type of the alert. Alert types include:
  - Replication
  - Storage Subsystem
  - Computer
  - Fabric
  - Switch
  - Port
  - Tape Library
- **Disk Manager Storage Subsystems table**
- **Fabric Manager Fabrics table**
- **Tape Manager Tape Libraries table**
- **Element Manager page** (of the Configuration Utility)



- **Topology Viewer objects** (see the following table)

| Topology Object        | Icon                                                                                |
|------------------------|-------------------------------------------------------------------------------------|
| Fabric                 |    |
| Switch                 |    |
| Switch port            |    |
| Tape library           |    |
| Storage subsystem      |    |
| Storage subsystem port |    |
| File system volume     |    |
| Storage disk           |   |
| Storage pool           |  |

## Launch menu items

The following are the possible menu items that can be displayed in the launch menu:

- **Launch in context defined application.** This menu item is displayed for each application that is defined by a launch descriptor file. Tivoli Storage Productivity Center provides launch descriptor files for the following applications:
  - DS8000 Storage Manager
  - The management console for the storage systems that are listed under **Management GUI**.
  - Tivoli Storage Productivity Center for Replication

If the corresponding launch descriptor files are found for the applications, the following menu items are displayed:

### DS8000 Storage Manager

Click this menu item to start DS8000 Storage Manager.

### Management GUI

Click this menu item to start the management console for the following storage systems:

- SAN Volume Controller 6.1 or later

- Storwize V7000
- Storwize V7000 Unified

### Replication Storage Systems Overview


Click this menu item to start Tivoli Storage Productivity Center for Replication. The window that is opened in the Tivoli Storage Productivity Center for Replication GUI depends on the context for the entity that you selected. For example, if you right-click a storage device on the **Element Manager** tab of the Configuration Utility and click **Replication Storage Systems Overview**, the Tivoli Storage Productivity Center for Replication Storage Systems window opens.

If the launch descriptor file cannot be found, the **Launch file items error** menu item is displayed. Click this item to view a description of the problem. The **Launch file items error** menu item is also displayed in situations such as the IBM Tivoli Integrated Portal server is not installed or is not accessible, the Tivoli Storage Productivity Center Data server is not accessible, or another system error occurred.

- **External Tools.** This menu item shows any external tools that are defined in the Configuration Utility. Examples of external tools include web-based user interfaces and local executable files. For example, you might add the address of a website that you frequently access or an application that you often use such as a text editor or a spreadsheet.
- **Element Manager.** This menu item shows any element managers that were discovered for the device by a Storage Management Initiative Specification (SMI-S) query. To avoid menu item duplication, this item is not displayed when the associated element manager is defined by a launch descriptor file. In this situation, the element manager is displayed as a separate menu item.
- **User Defined Tools.** This menu item shows any user-defined element managers that are defined in the Configuration Utility. These element managers differ from the element managers in the **Element Manager** menu because they are user-defined and are not discovered by a SMI-S query.

### Example launch menu

The following figure shows an example launch menu for a storage system in the **L0:Storage** topology view. In this example, the menu items **Replication Storage Systems Overview**, **DS8000 Storage Manager**, and **External Tools** are displayed. The **Element Manager** menu item is not displayed because no element managers other than **DS8000 Storage Manager** were discovered. The **User Defined Tools** menu item is not displayed because there are no user-defined element managers that are defined in the Configuration Utility.



## Using the single sign-on feature with launch or launch in context

*Single sign-on* is an authentication process in which you can access more than one system or application by entering a single user ID and password. Single sign-on enables you to start applications from Tivoli Storage Productivity Center without providing logon credentials for the application. Single sign-on is available only for Tivoli Storage Productivity Center for Replication, DS8000, SAN Volume Controller 6.1 and later, and Storwize V7000.

For Tivoli Storage Productivity Center for Replication Storage Manager, single sign-on is enabled by default. For DS8000 Storage Manager, SAN Volume Controller, and Storwize V7000, single sign-on must be enabled.

To enable single sign-on, see the following topics in the Tivoli Storage Productivity Center information center:

- Configuring Tivoli Storage Productivity Center for DS8000 LDAP authentication
- Configuring Tivoli Storage Productivity Center for SAN Volume Controller and Storwize V7000 single sign-on

### Related concepts:

“Single sign-on” on page 58

*Single sign-on* is an authentication process that enables you to enter one user ID and password to access multiple applications. For example, you can access IBM Tivoli Integrated Portal and then access IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication from Tivoli Integrated Portal using a single user ID and password. Single sign-on integrates with the launch and launch in context features to enable you to move smoothly from one application to another application.

---

## External tools

External tools include Web-based user interfaces and local executable files, scripts, and batch files that you want to access from IBM Tivoli Storage Productivity Center. For example, you might define the address of a Web site that you frequently access or an application that you often use (a text editor, a spreadsheet application, and so on) as an external tool. External tools such as executable files, scripts, and batch files must exist on the same system as the Tivoli Storage Productivity Center graphical user interface (GUI).

Both element managers and external tools are configured on the **Element Manager** tab of the Configuration Utility (**IBM Tivoli Storage Productivity Center > Configuration**). The following tasks can be performed for external tools in the **External Tools** section of the **Element Manager** tab. The tasks that you can perform depend on your Tivoli Storage Productivity Center authorization

- Add external tools (Administrator)
- Start external tools (all authorization levels)
- View external tools (all authorization levels)
- Edit external tools (Administrator)
- Remove external tools (Administrator)

## Adding external tools

Use the Configuration Utility to add external tools to IBM Tivoli Storage Productivity Center.

Possibilities for external tools include:

- Executable files, shell scripts, or batch files. These files must exist on the same system as the Tivoli Storage Productivity Center graphical user interface (GUI).
- Uniform Resource Locators (URLs) for search engines or web encyclopedias

To add an external tool in Tivoli Storage Productivity Center, follow these steps:

1. Expand the **IBM Tivoli Storage Productivity Center** node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** page.
2. In the **External Tools** section, select **Add Tool** from the **Select Action** list. The Add External Tool dialog box is displayed.
3. Define the information that is required to connect to the element manager in the following fields. Required fields are specified in bold in the dialog box. In the following field descriptions, required fields are marked by an asterisk.
4. The **Add External Tool** window opens. Enter definitions in the following fields:

#### **\*URL/Command**

Enter the URL or shell command representing the external tool. You can enter an IPv4 or IPv6 address depending on what is supported on the machine where the external tool is located:

- Enter an IPv4 address in the following format: ###.###.###.###. For example: 127.0.0.1
- Enter an IPv6 address. The preferred IPv6 address representation is written as eight groups of four hexadecimal digits: xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx, where each x is a hexadecimal digit representing 4 bits. For example: 2001:DB8:0000:1234:0000:0000:5678:ABCD. You can also specify IPv6 addresses using shortened formats that omit leading zeros (2001:DB8:0:1234:0:0:5678:ABCD) or use double colons (2001:DB8:0000:1234::5678:ABCD) in place of a series of zeros. You must include brackets [] around IPv6 addresses to separate those addresses from their port numbers. For example: [2001:DB8::1234:0000:0000:5678:ABCD]:9550.

**Note:** If the computer on which the element manager is located is IPv6-only, the Tivoli Storage Productivity Center servers must be installed on an IPv6-only or dual stack (IPv4 and IPv6 addresses are configured) computer.

**Label** Enter the name of the tool. The label name is shown in the **Name** field in the external tools table. If this field is blank, the **Name** field is blank.

**Description**

Enter a short description for the external tool.

**Save** Click this button to save the external tool. Information about the external tool is stored in the tools.xml file. This file is stored in the following locations:

- Windows: \Documents and Settings\Administrator\Application Data\IBM\TPC\config\
- UNIX and Linux: .TPC/config/

**Cancel**

Click this button to close the dialog box without saving the external tool.

## Modifying external tools

Use the Configuration Utility to modify information for external tools in IBM Tivoli Storage Productivity Center.

To modify an external tool in Tivoli Storage Productivity Center, follow these steps:

1. Expand the **IBM Tivoli Storage Productivity Center** node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **External Tools** section, select the tool that you want to modify from the external tools table.
3. Select **Modify Tool** from the **Select Action** list. The Modify External Tool dialog box is displayed.
4. Modify any of the following fields:

**URL/Command**

Enter the URL or shell command representing the external tool. You can enter an IPv4 or IPv6 address depending on what is supported on the machine where the external tool is located:

- Enter an IPv4 address in the following format: ###.###.###.###. For example: 127.0.0.1

- Enter an IPv6 address. The preferred IPv6 address representation is written as eight groups of four hexadecimal digits: `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx`, where each x is a hexadecimal digit representing 4 bits. For example: `2001:DB8:0000:1234:0000:0000:5678:ABCD`. You can also specify IPv6 addresses using shortened formats that omit leading zeros (`2001:DB8:0:1234:0:0:5678:ABCD`) or use double colons (`2001:DB8:0000:1234::5678:ABCD`) in place of a series of zeros. You must include brackets `[]` around IPv6 addresses to separate those addresses from their port numbers. For example: `[2001:DB8::1234:0000:0000:5678:ABCD]:9550`.

**Note:** If the computer on which the element manager is located is IPv6-only, the Tivoli Storage Productivity Center servers must be installed on an IPv6-only or dual stack (IPv4 and IPv6 addresses are configured) computer.

**Label** Enter the name of the tool. The label name is shown in the **Name** field in the external tools table. If this field is blank, the **Name** field is blank.

**Description**

Enter a short description for the external tool.

**Save** Click this button to save the external tool.

**Cancel**

Click this button to close the dialog box without saving the external tool.

## Starting external tools

Use the Configuration Utility to start external tools from IBM Tivoli Storage Productivity Center.

To start an external tool from Tivoli Storage Productivity Center, follow these steps:

1. Expand the **IBM Tivoli Storage Productivity Center** node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **External Tools** section, select the tool that you want to start from the external tools table.
3. Select **Launch Tool** from the **Select Action** list. The URL or command defined for the tool is run immediately.

## Removing external tools

Use the Configuration Utility to remove external tools from IBM Tivoli Storage Productivity Center.

To remove an external tool from Tivoli Storage Productivity Center, follow these steps:

1. Expand the **IBM Tivoli Storage Productivity Center** node in the navigation tree, click **Configuration Utility**, and open the **Element Manager** tab.
2. In the **External Tools** section, select the tool that you want to remove from the external tools table.
3. Select **Remove Tool** from the **Select Action** list. The Remove External Tool dialog box is displayed.
4. Click **OK**. The element manager is deleted from the external tools table.

To exit the window without removing the external tool, click **Cancel**.

**Note:** If the computer on which the element manager is located is IPv6-only, the Tivoli Storage Productivity Center servers must be installed on an IPv6-only or dual stack (IPv4 and IPv6 addresses are configured) computer.

---

## Working with invoices

With IBM Tivoli Storage Productivity Center you can develop invoice definitions and create invoices associated with data or resource usage.

Using Tivoli Storage Productivity Center you can:

- Charge users for the amount of *storage they are consuming* across a network and within relational database management system (RDBS) instances
- Charge users for the amount of *space they own* within the database and tablespace of RDBMS instances.

With Tivoli Storage Productivity Center you can collect usage information on a departmental, group, or user level. You can allocate costs by storage usage by user, disk capacity by computer, tablespace, or file system or physical device. You can create cost centers by creating user, computer, or tablespace groups, allowing organizations to chargeback individuals or business units for their storage usage. By understanding the costs associated with existing and future storage usage, you can improve the use of that storage and reduce the costs associated with its maintenances and upgrades.

Tivoli Storage Productivity Center uses the storage usage information gathered by probes and scans to generate invoices that chargeback for storage usage. Chargeback is the process of breaking down costs associated with data or resource usage and then distributing those costs back to the organizational consumers.

To collect the information needed for chargeback, you need to run probes and scans:

### Probes

Define probes when you want to charge by Disk Capacity. Probes collect statistics on the storage assets in your enterprise, such as computers, disk controllers, hard disks, file systems, storage subsystems, instances, logs, RDBMS objects, and so on. The results of probes are stored in the database repository and are used to supply the data necessary for Tivoli Storage Productivity Center's reporting function.

**Scans** Define scans when you want to charge by OS User groups. Scans collect statistics on the usage and trending of your actual storage consumption. The results of scans are stored in the database repository and are used to supply the data necessary for Tivoli Storage Productivity Center's Capacity, Usage, Usage Violations, and Backup Reporting facilities.

In addition to providing basic storage usage invoicing, Tivoli Storage Productivity Center integrates with the chargeback systems already implemented in your environment. It provides you with a higher level, application-specific CIMS output format. CIMS is IBM Tivoli Usage and Accounting Manager (ITUAM) which provides IT financial knowledge by measuring, analyzing, reporting, and billing based on usage and costs of shared computing resources. With Tivoli Storage Productivity Center, you can export chargeback data directly into ITUAM, giving you the ability to integrate your storage usage data with other enterprise chargeback information and processes.

## Planning for chargeback

Before generating data that will be used for chargeback, it is a good idea to plan ahead and gather basic information about your system and organization.

### Understand your environment

Understanding the users, data, and resources within your organization will help you determine the most reasonable and enforceable chargeback policies. In the planning stages of developing a chargeback policy, it will help you to understand such things as:

- The types of data stored within your environment
- Patterns of data usage
- Different types of users and user groups
- The IT resources that exist within your organization

### Determine the chargeback policy you want to implement

It is important to learn and understand the advantages and disadvantages of the charge back policy you want to implement. For example, with IBM Tivoli Storage Productivity Center, you can:

- Divide the cost of storage usage by all users
- Divide the cost of storage usage by individual users or groups of users
- Charge by resource
- Charge by individual usage of resources

### Define monitoring groups

Organize the IT resources within your environment into functional groups. Defining groups enables you to classify your data and resources into billable entities, which in turn provides the information on who is using what storage. This enables you to bill specific groups according to their usage of storage space. See "Creating a group" on page 371.

#### Computer groups

Computer groups collect operating system and RDBMS storage statistics on computers that are related to each other or belong to a common group such as a department, a location, a territory, and so forth. Examples of computer groups are Sales, Accounting, Development, IT.

A computer can belong to one computer group only.

#### User groups

User groups are users that are related to each other or belong to a common group such as a department, a location, a territory, and so forth. Examples of user groups are Administrators, Developers, Marketing staff, HR staff.

A user can belong to one user group only.

#### OS User groups

OS user groups are users that are grouped together within the operating system.

#### Database-tablespace groups

Database-tablespace groups monitor entire databases or tablespaces across multiple instances or databases. Examples of tablespace groups are System Tablespaces, System Databases, databases-tablespaces related to a specific application or applications, and so forth.



A database-tablespace can belong to one database-tablespace group only.

### **Determine what information you want to gather**

Determining what information you want to gather about the storage resources within your enterprise is critical to helping you implement a chargeback strategy. Probes and scans are the monitoring jobs you can define to collect information about operating system and RDBMS storage resources. The information gathered by these jobs is reflected in the data included within a Tivoli Storage Productivity Center invoice.

As you become more familiar with how a specific charge back model works for your organization, you can continue to review and implement to more advanced charge back models.

## **Setting up invoices**

This topic takes you through the task of setting up an invoice.

To set up an invoice, complete the following steps:

1. In the Navigation Tree pane, expand **Data Manager for Chargeback > Parameter Definition**. Data Manager for Chargeback - Configuration window opens.
2. In the **Output Format** area, specify the format for the output. You can click either of the following options:

#### **Invoice**

Creates a formatted text file suitable for opening in or importing into word processing applications.

**CIMS** Creates a text file with comma-separated fields suitable for importing into a ITUAM application.

3. In the **Type of Charges** area, specify the information that you want the invoice to contain. You can select one or more of the following check boxes:

#### **OS Storage Usage By User**

The resulting invoice includes the following information:

- Total space consumed by user groups and the cost of that usage
- Total space consumed by each user in the environment and the cost charged to each user

#### **OS Disk Capacity By Computer**

The resulting invoice includes the following information:

- Total space consumed by computer groups and the cost of that usage
- Total space consumed on each computer in the environment and the cost of that usage

#### **Storage Usage By Database User**


The resulting invoice includes the following information:

- Total space consumed by DB user groups and the cost of that usage
- Total space consumed by each DB user in the environment and the cost charged to each user

#### **Total Size By Database-Tablespace**

The resulting invoice includes the following information:

- Total space consumed by database-tablespace groups and the cost of that usage

- Total space consumed on each database or tablespace in the environment and the cost of that usage
4. (Invoice output format only) For each type of charge that you selected in step 3 on page 627, select the unit of measurement (kilobyte, megabyte, or gigabyte) and type a numerical value. For example, if you want to charge \$1.25 per megabyte of storage usage, select **Megabyte** and type 1.25 in the **Cost per** field. (CIMS output format only) If you select CIMS as the output type, the **Cost Per** field will be deselected. This is because you will set the rates and costing within the ITUAM application.
  5. If you selected **OS Storage Usage by User** or **Storage Usage by Database User**, select the profile that is used in the scan job. By default, this is either TPCUser.Summary By Owner for **OS Storage by User** or TPCUser.Db User Space for **Storage Usage by Database User**.
  6. Click  to save the invoice definition.

Invoice definitions are not saved in actual files, but the last invoice definition saved is the one that will be used to create the invoice.

## Generating an invoice

Generate an invoice and save it as a file.

To generate an invoice, complete the following steps:

1. In the Navigation Tree pane, expand the **Data Manager for Chargeback** node.
2. Click **Online Execution**. The Online Execution page opens.
3. Click the **Browse** button. The Open window opens, listing previously generated invoice files in the default directory. The default directory, or the directory where you have chosen to save the invoice files, is not part of the IBM Tivoli Storage Productivity Center application but on the computer where Tivoli Storage Productivity Center is installed.
  - a. Click an existing file or type a new file name, including extension, in the **File name** text field. A text file (.txt) is the default file type for both output types.
 

**Note:** Unless you save the invoice as a separate file name, the default file (StorageChargeback.txt) will be overwritten each time you run an invoice. It is recommended that you save your invoices as a separate file name each time you run them.
  - b. Click the **Open** button. You will be returned to the Online Execution page. Alternatively, in the Online Execution page, without clicking the **Browse** button, type the full path name of the invoice file to be generated in the **Output File** text field.
4. Click the **Generate Output File** button to generate the invoice.

The invoice is saved to the file name and path you specified.

## Viewing and printing invoices

This topic discusses viewing or printing a generated invoice in both Invoice and CIMS output formats.

If you have generated the invoice in the Invoice output format, you can view and print that invoice with any text editor or word processing program. If you chose CIMS as the output format, use your ITUAM application to work with the invoice data.

To view and print an invoice generated in the Invoice output format, complete the following steps:

1. Start the text editor or word processing program.
2. Follow the steps in the word processing program to open the invoice. Open the invoice from the location you specified in the **Output File** field in the Online Execution page.
3. View the file and make any changes. If you make any changes to the file, be sure and save it under a different name.
4. Print the file from the word processing program.

To view and print an invoice generated in the CIMS output format, integrate IBM Tivoli Storage Productivity Center accounting data into the Tivoli Usage and Accounting Manager, by following the instructions in the Tivoli Usage and Accounting Manager Information Center. Alternatively, you can view the original invoice data in a text editor or word processing program. Remember that the data will be in the form of comma-separated fields.

## Sample Invoices

This section gives examples of sample invoices.

### OS Storage Usage By User



OS Storage Usage By User

Cost per: MegaB... 0.05

Profile: TPCUser.Summary By Own...




Figure 49. Sample Invoice — OS Storage Usage By User

### OS Disk Capacity By Computer

OS Disk Capacity By Computer

Cost per:

IBM Tivoli Storage Resource Manager for Chargeback  
Computer Disk Space Invoice

| Name                                                                         | SPACE (MB) | Cost     |
|------------------------------------------------------------------------------|------------|----------|
| <i>The name of a user group being charged for storage usage.</i>             |            |          |
| companyABC Accounting                                                        |            | 0.05/MB  |
| ACC-server1                                                                  | 38,167     | 1,908.35 |
| ACC-server2                                                                  | 38,167     | 1,908.35 |
| group total                                                                  | 76,334     | 3,816.70 |
| <i>The total storage usage and cost of that storage usage for the group.</i> |            |          |
| companyABC.Development                                                       |            | 0.05/MB  |
| DEV-server1                                                                  | 17,367     | 868.35   |
| DEV-server2                                                                  | 19,474     | 973.70   |
| DEV-server3                                                                  | 30,521     | 1,526.05 |
| DEV-server4                                                                  | 26,056     | 1,302.80 |
| DEV-server5                                                                  | 19,093     | 954.65   |
| DEV-server6                                                                  | 28,616     | 1,430.80 |
| DEV-server7                                                                  | 9,542      | 477.10   |
| group total                                                                  | 150,669    | 7,533.45 |
| <i>Total storage space used by group.</i>                                    |            |          |
| <i>Total cost charged to group for storage usage.</i>                        |            |          |

Figure 50. Sample Invoice — OS Disk Capacity by Computer

### Storage Usage By Database User

Storage Usage By Database User

Cost per: MegaB... 0.185

Profile: TPCUser.Db User Spa...

DataBase User Usage Invoice

IBM.Default User Group — *The name of a user group being charged for storage usage.*

| Name         | SPACE (MB) | Cost  |
|--------------|------------|-------|
| APPI         | 35         | 6.47  |
| APP2         | 250        | 46.25 |
| SYSTEM       | 8          | 1.48  |
| group totals | 293        | 54.20 |
| run total    | 293 MB     | 54.20 |

*The names of the owners within the group.*

*This column displays the total space usage of a user.*

*The cost per megabyte of used space and the total cost for each owner/group.*

*The total storage usage and cost of that storage usage for the group.*

*Total storage space used by group.*

*Total cost charged to group for storage usage.*

Figure 51. Sample Invoice — Storage Usage by Database User

### Total Size By Database -Tablespace

Total Size By Database-TableSpace

Cost per:

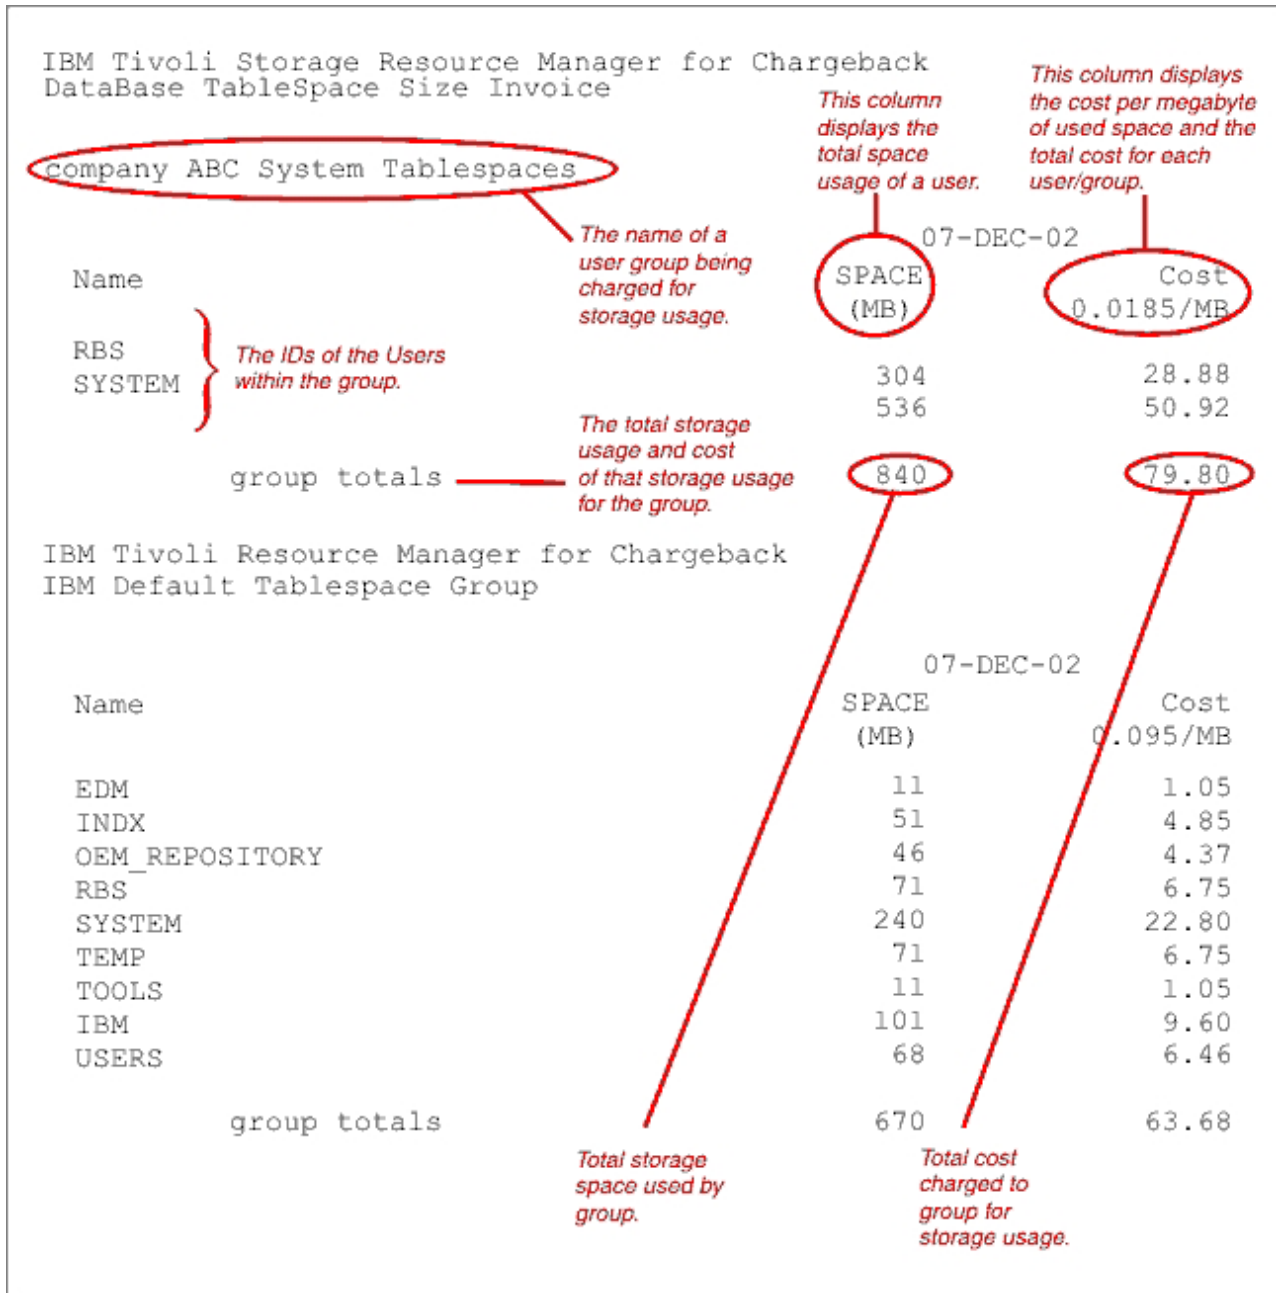


Figure 52. Sample Invoice — Total Size by Database-Tablespace





---

## Chapter 5. Managing replication

Manage replication components, including setting up the management servers, adding storage connections and paths, creating sessions and copy sets, setting up data replication, performing planned maintenance, and recovering from a disaster.

---

### Administering

Administer IBM Tivoli Storage Productivity Center for Replication to authorize users, start and use the graphical user interface (GUI), start and stop services, and many other administrative tasks.

### Starting and stopping IBM Tivoli Storage Productivity Center for Replication

Use these procedures to start and stop IBM Tivoli Storage Productivity Center for Replication, including the embedded IBM WebSphere Application Server and DB2.

#### Starting IBM Tivoli Storage Productivity Center for Replication

You start IBM Tivoli Storage Productivity Center for Replication by starting the embedded IBM WebSphere Application Server.

#### Starting IBM Tivoli Storage Productivity Center for Replication on Windows:

To start IBM Tivoli Storage Productivity Center for Replication on Windows, perform one of these procedures:

- From the desktop, perform these steps:
  1. Click **Start > Control Panel > Administrative Tools > Services**.
  2. Right-click IBM WebSphere Application Server, and then click **Start**.
- From the command line, enter the following command:

```
install_root/eWAS/profiles/CSM/bin/startServer.bat server1
```

#### Starting IBM Tivoli Storage Productivity Center for Replication on AIX and Linux:

To start IBM Tivoli Storage Productivity Center for Replication on AIX and Linux, enter the following command from the command line:

```
install_root/eWAS/profiles/CSM/bin/startServer.sh server1
```

#### Starting IBM Tivoli Storage Productivity Center for Replication on z/OS:

For IBM Tivoli Storage Productivity Center for Replication Basic Edition for System z or IBM Tivoli Storage Productivity Center for Replication for System z that uses IBM WebSphere Application Server OEM Edition for z/OS, see the *IBM WebSphere Application Server OEM Edition for z/OS Configuration Guide* for information about how to start IBM WebSphere Application Server OEM Edition for z/OS.

For IBM Tivoli Storage Productivity Center for Replication for System z that uses IBM WebSphere Application Server on z/OS, perform these steps to start IBM Tivoli Storage Productivity Center for Replication:

1. Ensure that the IBM WebSphere Application Server hierarchical file system (HFS), DB2 HFS, and IBM Tivoli Storage Productivity Center for Replication HFS are all mounted on the UNIX System Services (USS). The root installation

directory for IBM Tivoli Storage Productivity Center for Replication is  
-PathPrefix-/usr/lpp/Tivoli/RM. The root production directory is  
-PathPrefix-/var/lpp/Tivoli/RM.

2. Run the **TSO ISH** command to obtain the name of the USS mount point for the instance as the ENV parameter (for example, SYS7.SYS7.BBOS001). This parameter is typically found in the /WAS\_HOME/ directory. This parameter contains the path to several IBM WebSphere Application Server scripts.
3. Run the following command from the System Display and Search Facility (SDSF) panel to start the IBM WebSphere Application Server:  

```
/START BB06ACR,JOBNAME=BBOS001,ENV=ENV_parameter
```

The initiator is BB06ACR, and the job name is BBOS001. When you start a BB06ACR, the BBODMNB also starts.

## Stopping IBM Tivoli Storage Productivity Center for Replication

You stop IBM Tivoli Storage Productivity Center for Replication by stopping the embedded IBM WebSphere Application Server.

### Stopping IBM Tivoli Storage Productivity Center for Replication on Windows:

To stop IBM Tivoli Storage Productivity Center for Replication on Windows, perform one of these procedures:

- From the desktop, perform these steps:
  1. Click **Start > Control Panel > Administrative Tools > Services**.
  2. Right-click IBM WebSphere Application Server, and then click **Stop**.
- From the command line, enter the following command:

```
install_root/eWAS/profiles/CSM/bin/stopServer.bat server1
```

### Stopping IBM Tivoli Storage Productivity Center for Replication on AIX and Linux:

To stop IBM Tivoli Storage Productivity Center for Replication on AIX and Linux, enter the following command from the command line:

```
install_root/eWAS/profiles/CSM/bin/stopServer.sh -username user_name
-password password server1
```

Where *user\_name* and *password* are your IBM WebSphere Application Server user name and password.

### Stopping IBM Tivoli Storage Productivity Center for Replication on z/OS:

For IBM Tivoli Storage Productivity Center for Replication Basic Edition for System z or IBM Tivoli Storage Productivity Center for Replication for System z that uses IBM WebSphere Application Server OEM Edition for z/OS, see the *IBM WebSphere Application Server OEM Edition for z/OS Configuration Guide* for information about how to stop IBM WebSphere Application Server OEM Edition for z/OS.

For IBM Tivoli Storage Productivity Center for Replication for System z that uses IBM WebSphere Application Server on z/OS, perform these steps to stop IBM Tivoli Storage Productivity Center for Replication:

1. Run the following command from the System Display and Search Facility (SDSF) panel to stop the IBM WebSphere Application Server daemon process:  

```
/STOP BBOS001
```

**Attention:** For WebSphere Application Server OEM Edition for z/OS, use the following command:

```
/STOP WASOM1
```

2. Run the following command from the System Display and Search Facility (SDSF) panel to stop the daemon process:

```
/STOP BBODMNB
```

**Attention:** For WebSphere Application Server OEM Edition for z/OS, use the following command:

```
/STOP WAS0M1D
```

3. Run the **/D A.L** command to display the running processes and verify that the stop command completed successfully. If you see the BBOS001 or BBODMNB processes, run the **/CANCEL** command.

## Starting and stopping DB2

Use the DB2 Command Line Processor to start and stop DB2.

**Note:** If you are using the zero administration embedded repository (Apache Derby), you do not need to start or stop the repository.

To start DB2, run the **db2start** command from the DB2 Command Line Processor.

To stop DB2, run the **db2stop** command from the DB2 Command Line Processor.

**Important:** All connections to DB2 must be removed, or the **db2stop** command will fail. It is recommended that you stop the embedded IBM WebSphere Application Server, IBM Tivoli Storage Productivity Center for Replication, and all other applications accessing DB2 to ensure there are no connections to DB2 before attempting to stop DB2.

## Verifying that components are running

This information describes how to verify that the IBM Tivoli Storage Productivity Center for Replication components are running.

### Verifying that IBM WebSphere Application Server is running

This information describes how to verify that IBM WebSphere Application Server is running in your environment.

To verify that IBM WebSphere Application Server is running on Windows, click **Control Panel > Administrative Tools > Services** to view the Services panel. Find the IBM WebSphere Application Server V6.1 service and confirm that its Status is "Started".

Alternatively, locate the `server1.pid` file in the `install_root/eWAS/profiles/CSM/logs/server1` directory. This file contains the process identification (PID) number for the IBM WebSphere Application Server.

For a Windows system, use the Task Manager console to confirm that this PID is active by looking for the PID on the Processes panel. The Image Name is `java.exe`. On AIX or Linux, use the **ps** command to verify that the process is running.

### Verifying that the IBM Tivoli Storage Productivity Center for Replication server is running

This information describes how to verify whether the IBM Tivoli Storage Productivity Center for Replication server is running.

Perform one or more of these tasks to determine whether the IBM Tivoli Storage Productivity Center for Replication server is running:

- Start the IBM Tivoli Storage Productivity Center for Replication GUI or start up the IBM Tivoli Storage Productivity Center for Replication command line interface shell. If either of these methods is successful, then the IBM Tivoli Storage Productivity Center for Replication server is running.
- Determine whether IBM WebSphere Application Server is running. If IBM WebSphere Application Server is not running, then the IBM Tivoli Storage Productivity Center for Replication server is also not running.
- If IBM WebSphere Application Server is running, perform these steps to determine if IBM Tivoli Storage Productivity Center for Replication is configured correctly in IBM WebSphere Application Server:
  1. From a command prompt, change to the `install_root/eWAS/profiles/CSM/bin` directory.
  2. Enter the `wsadmin` command.
  3. Enter the `$AdminApp list` command.  
If you see CSM and CSMGUI in the resulting list, then IBM Tivoli Storage Productivity Center for Replication is configured correctly.
- View the `csmTrace.log` file in the `install_root/eWAS/profiles/CSM/logs` directory. If the `csmTrace.log` is being updated regularly and increasing in size, then IBM Tivoli Storage Productivity Center for Replication is running.

### Verifying that DB2 is running

This information describes how to verify that DB2 is running in your environment.

If IBM Tivoli Storage Productivity Center for Replication is configured to use DB2 and DB2 is not running, IBM Tivoli Storage Productivity Center for Replication cannot start correctly. If this is the case, you will find log entries similar to the following in the `csmTrace.log` files:

```
[2007-02-23 16:30:30.369-05:00] server.startup : 0 RepMgr D
com.ibm.csm.server.CSM databaseReady TRACE: The database is not ready for
connections. Will retry in 10 seconds
```

```
[2007-02-23 16:30:40.403-05:00] server.startup : 0 RepMgr D DBService
isDatabaseReady TRACE: [IBM][CLI Driver] SQL1032N No start database manager
command was issued. SQLSTATE=57019 DSR0010E: SQL State = 57019,
Error Code = -1,032
```

If DB2 is running and IBM Tivoli Storage Productivity Center for Replication can successfully connect to it, the following log entry can be found in the `csmTrace.log` file:

```
[2007-02-23 16:30:55.965-05:00] server.startup : 0 RepMgr D
com.ibm.csm.server.CSM databaseReady TRACE: The database is ready to go.
```

## Starting the IBM Tivoli Storage Productivity Center for Replication GUI

After IBM Tivoli Storage Productivity Center for Replication has been installed, you can access the IBM Tivoli Storage Productivity Center for Replication GUI.

Ensure that you used IBM Tivoli Storage Productivity Center Standard Edition to install Tivoli Integrated Portal and IBM Tivoli Storage Productivity Center for Replication.

Go to the following URL to start the IBM Tivoli Storage Productivity Center for Replication GUI. Note that the Web address is case sensitive.

```
https://host_name:port/CSM
```

**Note:**

- The default HTTP port is 3080, and the default HTTPS port is 3443.
- You can verify the ports that are correct for your installation in the `install_root/eWAS/profiles/profile_name/properties/portdef.props` file. The ports is defined by the `WC_defaulthost` (HTTP port) and `WC_defaulthost_secure` (HTTPS port) properties within the file.

To log in to the command line interface, use the CLI user ID and password that you entered during the installation of IBM Tivoli Storage Productivity Center for Replication. This ID and password are the same as the Administrator ID and password that you entered when you installed the product.

## Identifying the version of IBM Tivoli Storage Productivity Center for Replication

This topic tells how to view the IBM Tivoli Storage Productivity Center for Replication version.

The version of code on the IBM Tivoli Storage Productivity Center for Replication server determines the available features and enhancements added to the product. By clicking **About**, located in the Main menu, you can view the version and release of the IBM Tivoli Storage Productivity Center for Replication. The `ver` command described in the *IBM Tivoli Storage Productivity Center for Replication for System z Command-line Interface User's Guide*, also displays the current version of the product.

## Backing up and restoring IBM Tivoli Storage Productivity Center for Replication configuration data

You can back up the entire IBM Tivoli Storage Productivity Center for Replication database. You can also back up copy sets in a specific session. You can then use the backup files to restore a previous configuration or recover from a disaster.

### Back up and recovery

You can backup and recover copy set data for a specific session and the complete Tivoli Storage Productivity Center for Replication database.

### Copy sets

You can export data about all copy sets in a specific session to maintain a backup copy that you can use to recover if you lose the session or to upgrade to a different management server.

When exporting copy sets, Tivoli Storage Productivity Center for Replication takes a snapshot of the session at that point in time and saves the data in a comma separated value (CSV) file, which you can view or edit in a spreadsheet program such as Microsoft Excel. The exported CSV file includes the session name, session type, date that the data was exported, and the copy sets for each role pair. There is one line per copy set, and the volumes in the copy set are separated by a comma (for example: `ESS:2105.FCA57:VOL:17C7,ESS:2105.12043:VOL:17C7`).

The following example illustrates the content of the CSV file for a FlashCopy session. Note that the first valid row must contain the appropriate role names for the session. The order of the copy sets does not matter, and you can include extra roles. A copy set is created from each row that follows the role names. All rows

must have data in each column to be a valid row. Note that the number sign (#) indicates that the line is a comment. Lines that are comments are ignored.

```
#Session1,
#FlashCopy,
#Oct 2, 2009 10:03:18 AM
```

```
H1,T1
DS8000:2107.FRLL1:VOL:1004,DS8000:2107.FRLL1:VOL:1104
DS8000:2107.FRLL1:VOL:1011,DS8000:2107.FRLL1:VOL:1101
DS8000:2107.FRLL1:VOL:1005,DS8000:2107.FRLL1:VOL:1105
```

**Important:** You must manually save this file on the local system when you export copy sets from the Tivoli Storage Productivity Center for Replication Web interface.

## IBM Tivoli Storage Productivity Center for Replication database

Tivoli Storage Productivity Center for Replication database contains all product data, including data about storage systems, sessions, copy sets, paths, user administration and management servers. You can back up this data and use the backup file to recover from a disaster or restore a previous configuration.

**Important:** You must have Administrator privileges to back up and recover the database.

The current data is stored in a new file each time you create a backup. The backup file is named *yyyyMMdd\_HHmmsSSS.zip*, where *yyyy* is the year, *MM* is the month, *dd* is the day, *HH* is the hour, *mm* is the minute, *ss* is the seconds, *SSS* is the milliseconds when the backup command was run. It is your responsibility to delete backup versions that are no longer needed.

By default, the backup file is stored in the *install\_root/eWAS/profiles/CSM/database/backup* directory. You can change the default location by editing the **db.backup.location** property in *rmserver.properties* file, which is located in the *websphere\_home/AppServer/profiles/websphere\_profile/properties* directory.

The backup file contains the Tivoli Storage Productivity Center for Replication database data at the time the backup was performed. Any changes that were made after the backup are not reflected when the backup files are used to restore an Tivoli Storage Productivity Center for Replication database. It is recommended that you create a new backup file:

- After changing the Tivoli Storage Productivity Center for Replication database data, such as adding or deleting a storage system, changing properties, and changing user privileges
- After an Tivoli Storage Productivity Center for Replication session changes direction. For example, if a Metro Mirror session was copying data from H1 to H2 when the backup was taken, and later, the session was started in the H2 to H1 direction. The session must be in the Prepared state before you create the backup.
- After a site switch has been declared and the Enable Copy To Site command has been issued. After you create a backup, consider deleting the previous backup to prevent Tivoli Storage Productivity Center for Replication from starting the copy in the wrong direction.

When you create a backup, ensure that all Tivoli Storage Productivity Center for Replication sessions are either in the Defined, Prepared or Target Available state.

Restoring the Tivoli Storage Productivity Center for Replication database from a backup copy puts Tivoli Storage Productivity Center for Replication back to the point in time when the backup was made. Relationships that exist on the storage systems that were created by Tivoli Storage Productivity Center for Replication after the backup was made will no longer be managed by Tivoli Storage Productivity Center for Replication until you add the copy set to the session and Tivoli Storage Productivity Center for Replication assimilates the relationship into the session. Copy sets that were deleted after the backup will be restored and a subsequent Start command to the session will create new relationships; therefore, you must remove the deprecated copy sets before issuing the Start command.

After restoring a Global Mirror session, you must stop the Global Mirror master and subordinates before restarting the Global Mirror session. Refer to your DS6000 and DS8000 storage system documentation for more information.

## Backing up the Tivoli Storage Productivity Center for Replication database

This topic describes how to create a backup of the IBM Tivoli Storage Productivity Center for Replication database, including data about storage systems, sessions, copy sets, user administration and management server configuration.

To back up the Tivoli Storage Productivity Center for Replication database, run the **mkbackup** command from the command line, for example:

```
csmdi> mkbackup
```

### Prerequisites:

- You must have Administrator privileges to run this command.
- This procedure applies to only the zero-administration embedded repository. This procedure *is not* applicable when DB2 is being used as the persistent datastore for the IBM Tivoli Storage Productivity Center for Replication database. For information about restoring your DB2 environment, refer to your DB2 documentation.

By default, the backup file is stored in the *install\_root/eWAS/profiles/CSM/database/backup* directory. You can change the default location by editing the **db.backup.location** property in *rmserver.properties* file, which is located in the *websphere\_home/AppServer/profiles/websphere\_profile/properties* directory.

## Restoring the IBM Tivoli Storage Productivity Center for Replication database

You can restore IBM Tivoli Storage Productivity Center for Replication database that was previously backed up to the local system.

### Important:

- Restoring the database does not require Administrator privileges. However, you must be able to access the files on the IBM Tivoli Storage Productivity Center for Replication server that are listed in the procedure.
- This procedure applies to only the zero-administration embedded repository. This procedure *is not* applicable when DB2 is being used as the persistent datastore for the IBM Tivoli Storage Productivity Center for Replication database. For information about restoring your DB2 environment, refer to your DB2 documentation.

Perform these steps to restore the IBM Tivoli Storage Productivity Center for Replication database from a backed up version:

1. Stop IBM Tivoli Storage Productivity Center for Replication on the active management server by running the **stopServer** command from a command line.
2. Delete the *install\_root/eWAS/profiles/CSM/database/csmdb* directory and all contents in it.
3. Uncompress the backup file into the *install\_root/eWAS/profiles/CSM/database* directory.
4. If IBM Tivoli Storage Productivity Center for Replication is running on z/OS, change the permissions of the csmdb directory by running the following commands:
  - `chgrp -R $WAS_GROUP csmdb`
  - `chmod -R u+rwx csmdb`
  - `chmod -R g+rwx csmdb`
  - `chmod -R o+r csmdb`
5. Restart IBM Tivoli Storage Productivity Center for Replication on the active management server by running the **startServer** command from a command line.
6. Resolve any changes that might have occurred since the backup was created.
7. Start the IBM Tivoli Storage Productivity Center for Replication sessions using the appropriate start commands. The start commands reestablish the relationship between the volume pairs and synchronize data on those volumes.
8. If you have a standby management server, reestablish that standby relationship to update the database on the standby server.

### Exporting copy set data

You can export data about all copy sets in a specific session, to maintain a backup copy that can be used to recover if you lose your session or upgrade to a different server.

Perform these steps to export the copy sets in a specific session:

1. In the navigation tree, select **Sessions**. The Sessions panel is displayed.
2. Select the session for which you want to export copy sets.
3. Select **Export Copy Sets** from the **Actions** list, and click **Go**. The Export Copy Set wizard displays the status of the export and a link to the exported file.
4. Click that link and save the file to the local system.

**Important:** You must save the file to your local system. After you close the panel, the data will be lost.

5. Click **Finish**.

### Importing copy set data

You can import copy set data that was previously exported to a comma separated value (CSV) file.

Perform the following steps to import copy sets into an existing session:

1. In the navigation tree, select **Sessions**. The Session panel is displayed.
2. Select the session for which you want to import copy sets.
3. Select **Add Copy Sets** from the **Actions** list, and click **Go**. The Add Copy Sets wizard is displayed.
4. Select **Use a CSV file to import copy sets**.
5. Type the location and name of the CSV file to import, or use **Browse** to select the file. Then, click **Next**.



6. Verify that the matching results were successful, and then click **Next**.
7. Select the copy sets you want to add, and then click **Next**.
8. Confirm the number of copy sets that you want to create, and click **Next**. A progress bar displays.
9. Click **Next**.
10. Verify the matches, and click **Finish**.

---

## Managing management servers

This section provides information about how to set up active and standby management servers, restore a lost connection between the management servers, or perform a takeover on the standby management server.

### Management servers

The *management server* is a system that has IBM Tivoli Storage Productivity Center for Replication installed. The management server provides a central point of control for managing data replication.

You can create a high-availability environment by setting up a *standby management server*. A standby management server is a second instance of Tivoli Storage Productivity Center for Replication that runs on a different physical system, but is continuously synchronized with the primary (or active) Tivoli Storage Productivity Center for Replication server. The active management server issues commands and processes events, while the standby management server records the changes to the active server. As a result, the standby management server contains identical data to the active management server and can take over and run the environment without any loss of data. If the active management server fails, you can issue the Takeover command to make the standby management server take over.

### Connecting the active management server to the standby management server

Ensure that the active management server is connected to the standby management server. This connection creates the *management server relationship* that begins the synchronization process. Each management server can be in only one management server relationship.

A management server relationship might become disconnected for a number of reasons, including a connectivity problem or a problem with the alternate server. Issue the Reconnect command to restore synchronization.

### Performing a takeover on the standby management server

If you must perform a takeover and use the standby server, ensure that you shut down the active management server first. You must ensure that you do not have two active management servers. If there are two active management servers and a condition occurs on the storage systems, both management servers respond to the same conditions, which might lead to unexpected behavior.

If you perform an action on the active management server when the servers are disconnected, the servers will be out of synch.

## Viewing the status of the management servers

You can view the status of the active and standby management servers from the Management Servers panel in the Tivoli Storage Productivity Center for Replication graphical user interface (GUI). If you are logged on to the active management server, the icons on this panel show the status of the standby management server. If you are logged on to the standby management server, the icons on this panel show the status of the active management server.

When the status is Synchronized, the standby management server contains the same data that the active management server contains. Any update to the active management server database is replicated to the standby server database.

## Managing volumes on storage systems

When you add direct connections, Hardware Management Console (HMC) connections, or z/OS connections on the active management server, Tivoli Storage Productivity Center for Replication automatically enables the management of attached extended count key data (ECKD™) volumes, non-attached count key data (CKD) volumes, and all fixed-block volumes on the storage system. To disable management of volumes on the storage system, use the volume protection function.

## Information specific to management servers in z/OS environments

If the standby management server is not in the active server z/OS sysplex, the standby server is not able to communicate with the storage systems using a z/OS connection; therefore, another connection must be made using a TCP/IP connection.

If DB2 is configured for data sharing mode across the z/OS sysplex, one of the Tivoli Storage Productivity Center for Replication servers must be configured to use the zero-administration embedded repository. If the embedded repository is not used, the two servers will overwrite the same data in the Tivoli Storage Productivity Center for Replication database.

## Ports

Tivoli Storage Productivity Center for Replication uses ports for communication with the management servers in a high-availability relationship, graphical user interface (GUI), command-line interface (CLI), and storage systems.

### Web browser ports

To launch the Tivoli Storage Productivity Center for Replication GUI, use one of these default ports: HTTP port 3080 or HTTPS port 3443.

If you upgraded from Tivoli Storage Productivity Center for Replication version 3.3 or earlier, use one of these default ports: HTTP port 9080 or HTTPS port 9443.

You can verify the ports that are correct for your installation in the *install\_root/ewAS/profiles/profile\_name/properties/portdef.props* file. The ports are defined by the WC\_defaulthost (HTTP port) and WC\_defaulthost\_secure (HTTPS port) properties within the file.

## Standby management server port

Tivoli Storage Productivity Center for Replication uses the default port 5120 for communication between the active and standby management server. This port number is initially set at installation time.

**Important:** The standby management server port number must be the same on both management servers in a high-availability relationship. If you change the port number on one management server, you must also change it on the other.

You can view the current port for each management server by clicking **Management Servers** in the navigation tree or from the **Health Overview** panel in the GUI or using the `lshaservers` command from the command line interface.

## Client port

IBM Tivoli Storage Productivity Center for Replication client uses the default port 5110 to communicate with the graphical user interface and command line interface from a remote system. This port number is initially set at installation time.

**Important:** The client port number must be the same on both management servers in a high-availability relationship. If you change the port number on one management server, you must also change it on the other.

You can view the client port number on the local management server by clicking **About** in the navigation tree in the GUI or using the `whoami` command from the command line interface.

## Storage system ports

The following table lists the default ports for each storage type.

*Table 29. Storage system default ports*

| Storage System                                                                                                                                                     | Connection Type                        | Port       |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------|------------|
| <ul style="list-style-type: none"><li>• TotalStorage Enterprise Storage Server Model 800</li><li>• System Storage DS8000</li><li>• System Storage DS6000</li></ul> | Direct Connection                      | 2433       |
| <ul style="list-style-type: none"><li>• System Storage DS8000</li></ul>                                                                                            | Hardware Management Console Connection | 1750       |
| <ul style="list-style-type: none"><li>• System Storage SAN Volume Controller</li><li>• Storwize V7000</li><li>• Storwize V7000 Unified</li></ul>                   | Direct Connection                      | 443 and 22 |
| <ul style="list-style-type: none"><li>• The XIV system</li></ul>                                                                                                   | Direct Connection                      | 7778       |

Ensure that your network configuration is set up so that Tivoli Storage Productivity Center can send outgoing TCP/IP packets to the storage controllers. It is possible when adding the storage controllers to Tivoli Storage Productivity Center to set a specific port number for your storage controller.

Because there are typically multiple applications running on the management server, it is possible that port conflicts might arise if other applications attempt to

use the same ports that IBM Tivoli Storage Productivity Center for Replication is using. Use the **netstat** command to verify which ports the various applications on the management server are using.

When you add a storage system to the Tivoli Storage Productivity Center for Replication configuration, the port field is automatically populated with the appropriate value. If you want to use different ports, you can change them by clicking **Storage Systems** located in the navigation tree, clicking the storage system that you want to change, and then changing the port value in the View/Modify Details panel.

**Note:** The storage system must not be in a Connected state if you want to change port values.

If firewalls are being used in your configuration, ensure that none of these ports are being blocked. Also ensure that not only is the Tivoli Storage Productivity Center for Replication server granted access to reach the other components, but that the other components are granted access to reach the Tivoli Storage Productivity Center for Replication server.

**Note:** If the management server is running Windows Server 2008, configure the inbound and outbound rules for Tivoli Storage Productivity Center. To create a new outbound rule, launch the New Outbound Rule Wizard from the Windows Firewall with Advanced Security menu.

## SNMP alerts

This topic describes the SNMP alerts that are sent by IBM Tivoli Storage Productivity Center for Replication and the associated object IDs (OIDs).

SNMP alerts are sent during the following general events:

- Session state change
- Configuration change
- Suspending-event notification
- Communication failure
- Management Server state change

### Session state change SNMP trap descriptions

This topic lists the SNMP traps that are sent during a session state change. A different trap is sent for each state change. These alerts are sent by only the active management server.

A session state change SNMP trap is sent each time the session changes to one of the following states:

- Defined
- Preparing
- Prepared
- Suspended
- Recovering
- Flashing
- TargetAvailable
- Suspending
- (Metro Global Mirror only) SuspendedH2H3
- (Metro Global Mirror only) SuspendedH1H3

Table 30. Session state change traps

| Object ID (OID)          | Description                                                  |
|--------------------------|--------------------------------------------------------------|
| 1.3.6.1.4.1.2.6.208.0.1  | The state of session X has transitioned to Defined.          |
| 1.3.6.1.4.1.2.6.208.0.2  | The state of session X has transitioned to Preparing.        |
| 1.3.6.1.4.1.2.6.208.0.3  | The state of session X has transitioned to Prepared.         |
| 1.3.6.1.4.1.2.6.208.0.4  | The state of session X has transitioned to Suspended.        |
| 1.3.6.1.4.1.2.6.208.0.5  | The state of session X has transitioned to Recovering.       |
| 1.3.6.1.4.1.2.6.208.0.6  | The state of session X has transitioned to Target Available. |
| 1.3.6.1.4.1.2.6.208.0.19 | The state of session X has transitioned to Suspending.       |
| 1.3.6.1.4.1.2.6.208.0.20 | The state of session X has transitioned to SuspendedH2H3.    |
| 1.3.6.1.4.1.2.6.208.0.21 | The state of session X has transitioned to SuspendedH1H3.    |
| 1.3.6.1.4.1.2.6.208.0.22 | The state of session X has transitioned to Flashing.         |
| 1.3.6.1.4.1.2.6.208.0.23 | The state of session X has transitioned to Terminating.      |

## Configuration change SNMP trap descriptions

This topic lists the SNMP traps that are sent when the configuration changes. These alerts are sent by only the active management server.

Configuration change SNMP traps are sent after the following configurations changes are made:

- One or more copy sets have been added or deleted from a session  
 An alert is sent for each set of copy sets added to or removed from a session. Note that an alert for copy set changes is sent only once within 15 minutes of a configuration change, so you might not see alerts from successive changes that occur within that 15-minute period. For example, if you make a copy set configuration change that causes an alert to be sent at 10:41:01, and if you were to make additional copy set changes at 10:42:04 and 10:50:09, no alerts would be sent for these two changes because they occurred within the 15-minute minimum interval from the first alert.
- PPRC path definitions have been changed  
 An alert is sent for each path configuration change made.

Table 31. Configuration change traps

| Object ID (OID)         | Description                                                                                                                                      |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| 1.3.6.1.4.1.2.6.208.0.7 | One or more copy sets have been added or deleted from this session.<br><b>Note:</b> An event is sent for each session at least every 15 minutes. |
| 1.3.6.1.4.1.2.6.208.0.8 | Peer-to-Peer Remote Copy (PPRC) path definitions have been changed. An event is sent for each path configuration change.                         |

## Suspending-event notification SNMP trap descriptions

These SNMP traps that are sent during a suspending-event notification. The traps are sent by the active and standby management server.

Suspending-event notification SNMP traps indicate that a session has transitioned to a Severe status due to an unexpected error.

Table 32. Suspending-event notification traps

| Object ID (OID)         | Description                                                  |
|-------------------------|--------------------------------------------------------------|
| 1.3.6.1.4.1.2.6.208.0.9 | The session is in a Severe state due to an unexpected error. |

### Communication-failure SNMP trap descriptions

This topic lists the SNMP traps that are sent during a communication-failure. These alerts are sent by both the active and standby management servers.

Communication-failure SNMP traps are sent after the following events occur:

- A server times out attempting to communicate with a storage system.
- A server encounters errors attempting to communicate with a storage system.
- An active server terminates communication with a standby server as a result of communication errors.
- A standby encounters communication errors with an active server.

After an SNMP trap for a given failure is sent, it is not resent unless communication has been reestablished and failed again.

Table 33. Communication-failure traps

| Object ID (OID)          | Description                                                                                             |
|--------------------------|---------------------------------------------------------------------------------------------------------|
| 1.3.6.1.4.1.2.6.208.0.10 | Server X has timed out attempting to communicate with storage system Y.                                 |
| 1.3.6.1.4.1.2.6.208.0.11 | Server X has encountered errors attempting to communicate with storage system Y.                        |
| 1.3.6.1.4.1.2.6.208.0.12 | Active server X has terminated communication with standby server Y as a result of communication errors. |
| 1.3.6.1.4.1.2.6.208.0.13 | Standby server X has encountered communication errors with active server Y.                             |

### Management Servers state-change SNMP trap descriptions

This topic lists the SNMP traps that are sent when the state of the management server changes. These alerts are sent by both the active and standby management servers.

A management server state change SNMP trap is sent each time the management server changes to one of the following states:

- Unknown
- Synchronization Pending
- Synchronized
- Disconnected Consistent
- Disconnected

Table 34. Management Servers state-change traps

| Object ID (OID)          | Description                                                                                                                                            |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1.3.6.1.4.1.2.6.208.0.14 | The IBM Tivoli Storage Productivity Center for Replication Server Management Server connection X->Y has changed state to Unknown (previously Offline). |
| 1.3.6.1.4.1.2.6.208.0.15 | The IBM Tivoli Storage Productivity Center for ReplicationServer Management Server connection X->Y has changed state to Synchronized.                  |

Table 34. Management Servers state-change traps (continued)

| Object ID (OID)          | Description                                                                                                                                                                       |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1.3.6.1.4.1.2.6.208.0.16 | The IBM Tivoli Storage Productivity Center for Replication Server Management Server connection X->Y has changed state to Disconnected Consistent (previously Consistent Offline). |
| 1.3.6.1.4.1.2.6.208.0.17 | The IBM Tivoli Storage Productivity Center for Replication Server Management Server connection X->Y has changed state to Synchronization Pending.                                 |
| 1.3.6.1.4.1.2.6.208.0.18 | The IBM Tivoli Storage Productivity Center for Replication Server Management Server connection X->Y has changed state to Disconnected.                                            |

## Setting up a standby management server

You can set up a standby management server in two ways: setting up the management server you are logged in to as the standby management server, or designating another server as the standby management server.

**Note:** When you define a standby management server, the IBM Tivoli Storage Productivity Center for Replication code must be at the same level on both the standby and active management servers.

### Setting the local management server as the standby server

This topic describes how to set the management server on which you are currently logged in as the standby management server.

**Attention:** When you set the local management server, the server you are logged in at will be wiped of all session information, and replaced with the session information belonging to the server you specified.

**Important:** The standby management server port number must be the same on both management servers in a high-availability relationship. If you change the port number on one management server, you must also change it on the other.

Perform these steps to set the local management server as the standby server:

1. In the navigation tree, select **Management Servers**.
2. From the **Action** menu, select **Set This Server as Standby**, and click **Go**. This opens the Set This Server As Standby panel.
3. Enter the domain name or IP address of the desired active management server.
4. Click **OK** to connect to the active server. You have now designated the server you are logged in at as the standby server.

### Setting a remote management server as the standby server

This topic describes how to set up a management server on which you are not logged in as the standby server.

**Attention:** When you set a remote management server as the standby server, the remote management server is wiped of all session information and replaced with the session information belonging to the management server you specified.

**Important:** The standby management server port number must be the same on both management servers in a high-availability relationship. If you change the port number on one management server, you must also change it on the other.

Perform these steps to set a remote management server as the standby server:

1. In the navigation tree, select **Management Servers**.
2. Select **Define Standby** from the drop-down menu, and click **Go**. The Define Standby panel appears.
3. Type the domain name or IP address of the server that you are defining as the new standby management server. Log in to the standby management server by entering the username and password.
4. Click **OK** to connect to the standby management server.

## Reinstalling the primary server during an active session

This topic provides the steps to reinstall the primary IBM Tivoli Storage Productivity Center for Replication server if it develops a problem during an active session.

If there is an active session running when you need to reinstall the primary server, follow the steps listed here to reinstall the server without affecting the active session.

1. If the heartbeat is enabled, disable it:
  - a. In the navigation tree, select **Advanced Tools**.
  - b. Click **Disable Heartbeat**.
2. On the standby server, Server 2, issue a take-over. This makes Server 2 the active server. It is possible that the original active server, Server 1, is still listed on the Server 2 Management Servers page. If so, select **Remove Standby**.
3. Disable the heartbeat on Server 2, in case there are any problems.
4. Uninstall IBM Tivoli Storage Productivity Center for Replication on Server 1.
5. Reinstall IBM Tivoli Storage Productivity Center for Replication on Server 1.

**Note:** If no changes have been made to the configuration while Server 1 was being reinstalled, steps 6 and 7 are not necessary.

6. When IBM Tivoli Storage Productivity Center for Replication is running on Server 1, log into Server 2 and set Server 1 as the standby server for Server 2. This step copies the configuration from Server 2 to Server 1. This process takes a few minutes.
7. When the management servers status is synchronized, issue a take-over on Server 1. This makes Server 1 an active server, able to control sessions.

**Note:** It is possible that Server 2 is still listed on the Server 1 Management Servers page. If this is the case, select **Remove Standby**.

8. Disable the heartbeat on Server 1 to make sure this active server does not have any problems.

**Note:** If you do not need to reinstall IBM Tivoli Storage Productivity Center for Replication on Server 2, skip steps 9 and 10.

9. Uninstall IBM Tivoli Storage Productivity Center for Replication from Server 2.
10. Reinstall IBM Tivoli Storage Productivity Center for Replication on Server 2.
11. On Server 2, go to the Management Servers page and select the **Set This Server As Standby** option, entering the information for Server 1. When this step is complete, Server 1 is the active server, and Server 2 is the standby server.



12. When you are confident that the active server is running without any problems, enable the heartbeat again, if desired.

## Reconnecting the active and standby management servers

If the active and standby management servers become disconnected, reestablish that connection.

Perform these steps to cause the standby management server to become the active management server:

1. In the navigation tree, select **Management Servers**. The Management Servers panel is displayed.
2. Select **Reconnect** from the **Actions** list, and click **Go**.

## Performing a takeover on the standby management server

If the active management server fails, you can force the standby management server to take over monitoring and managing replication responsibilities.

**Important:** If the current active management server is still active, you must not attempt to control the replication environment simultaneously from both management servers. Instead, either reconfigure the current active management server to be a standby management server, or shut it down.

Perform these steps to cause the standby management server to become the active management server:

1. If the active management server is functioning, take it offline so you do not have two active management servers managing the same sessions.
2. Log in to the IBM Tivoli Storage Productivity Center for Replication Web interface running on the standby management server.
3. In the navigation tree, select **Management Servers**. The Management Servers panel is displayed.
4. Select **Takeover** from the **Actions** list, and click **Go**.
5. To reestablish high-availability, perform one of these steps:
  - Choose another server to be the standby management server. See instructions for setting up a standby management server.
  - Bring the failed management server back online, and then make that server the standby management server. See instructions for setting up a standby management server.
  - Bring the failed management server back online, and then make that server the active management server to return to the original configuration. Repeat the steps in this section and then add the original standby server as the standby server.

**Important:** Do *not* use the **Reconnect** command if you perform a takeover. You would use the **Reconnect** command when the active server loses its connection with the standby server; it reconnects the two servers. Do *not* use the **Reconnect** command after a takeover to reconnect to the original active server.

## Configuring SNMP

The SNMP community name has a default value of `public`.

To change the community name, modify or add the **csn.server.snmp\_community\_string** property in the `rmserver.properties` file, which is located in the `websphere_home/AppServer/profiles/websphere_profile/properties` directory.

## Adding SNMP managers

Use the **mksnmp** command to add an SNMP manager to the list of servers to which IBM Tivoli Storage Productivity Center for Replication sends SNMP alerts.

IBM Tivoli Storage Productivity Center for Replication uses management information base (MIB) files to provide a textual description of each SNMP alert sent by IBM Tivoli Storage Productivity Center for Replication. You must configure the SNMP manager to use both the `SYSAPPL-MIB.mib` and `ibm-TPC-Replication.mib` files. These MIB files are located in the `install_root\Scripts` directory. Follow the directions provided by your SNMP manager application to configure it to use the MIB files.

**Tip:** You can also find the MIB files on the installation CD in the `TPCRM/CSM-Client/etc` directory.

**Note:** IBM Tivoli Storage Productivity Center for Replication sends all SNMP alerts to each registered SNMP manager. SNMP alerts are not specific to any particular session, and all alerts for any session are sent. You cannot choose to send a subset of SNMP alerts.

## Changing the standby management server port number

The standby management server port is used for communication between the active and standby management server. This port is initially defined during the installation. You can manually change this port after installation.

**Important:** The standby management server port number must be the same on both management servers in a high-availability relationship. If you change the port number on one management server, you must also change it on the other.

1. Open the `rmserver.properties` file in the `websphere_home/AppServer/profiles/websphere_profile/properties` directory.
2. Modify the port number for the following property:  
`communications.haPort=5120`
3. Restart IBM Tivoli Storage Productivity Center for Replication. You must restart IBM Tivoli Storage Productivity Center for Replication to activate property changes. Properties are not synchronized between the IBM Tivoli Storage Productivity Center for Replication management servers and must be done on each IBM Tivoli Storage Productivity Center for Replication management server.

## Changing the client port number

The client port is used to log in to the graphical user interface and command line interface from a remote system. This port is initially defined during the installation. You can manually change this port after installation.

**Important:** The client port number must be the same on both management servers in a high-availability relationship. If you change the port number on one management server, you must also change it on the other.

1. Open the `rmservers.properties` file in the `websphere_home/AppServer/profiles/websphere_profile/properties` directory.
2. Modify the port number for the following property:  
`communications.port=5110`
3. Open the `repcli.properties` file in the `install_root/CLI` directory.
4. Modify the port number for the following property:  
`port=5110`
5. Restart IBM Tivoli Storage Productivity Center for Replication. You must restart IBM Tivoli Storage Productivity Center for Replication to activate property changes. Properties are not synchronized between the IBM Tivoli Storage Productivity Center for Replication management servers and must be maintained on each IBM Tivoli Storage Productivity Center for Replication management server.

---

## Managing storage systems

To replicate data among storage systems using IBM Tivoli Storage Productivity Center for Replication, you must add connections to the storage systems. After a storage system is added, you can associate a location, modify connection properties, set volume protection, and refresh the storage configuration for that storage system.

### Ports

Tivoli Storage Productivity Center for Replication use ports for communication with the management servers in a high-availability relationship, graphical user interface (GUI), command-line interface (CLI), and storage systems.

#### Web browser ports

To launch the Tivoli Storage Productivity Center for Replication GUI, use one of these default ports: HTTP port 3080 or HTTPS port 3443.

If you upgraded from Tivoli Storage Productivity Center for Replication version 3.3 or earlier, use one of these default ports: HTTP port 9080 or HTTPS port 9443.

You can verify the ports that are correct for your installation in the `install_root/eWAS/profiles/profile_name/properties/portdef.props` file. The ports is defined by the `WC_defaulthost` (HTTP port) and `WC_defaulthost_secure` (HTTPS port) properties within the file.

#### Standby management server port

Tivoli Storage Productivity Center for Replication uses the default port 5120 for communication between the active and standby management server. This port number is initially set at installation time.

**Important:** The standby management server port number must be the same on both management servers in a high-availability relationship. If you change the port number on one management server, you must also change it on the other.

You can view the current port for each management server by clicking **Management Servers** in the navigation tree or from the **Health Overview** panel in the GUI or using the `lshaservers` command from the command line interface.

## Client port

IBM Tivoli Storage Productivity Center for Replication client uses the default port 5110 to communicate with the graphical user interface and command line interface from a remote system. This port number is initially set at installation time.

**Important:** The client port number must be the same on both management servers in a high-availability relationship. If you change the port number on one management server, you must also change it on the other.

You can view the client port number on the local management server by clicking **About** in the navigation tree in the GUI or using the **whoami** command from the command line interface.

## Storage system ports

The following table lists the default ports for each storage type.

*Table 35. Storage system default ports*

| Storage System                                                                                                                                                     | Connection Type                        | Port       |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------|------------|
| <ul style="list-style-type: none"><li>• TotalStorage Enterprise Storage Server Model 800</li><li>• System Storage DS8000</li><li>• System Storage DS6000</li></ul> | Direct Connection                      | 2433       |
| <ul style="list-style-type: none"><li>• System Storage DS8000</li></ul>                                                                                            | Hardware Management Console Connection | 1750       |
| <ul style="list-style-type: none"><li>• System Storage SAN Volume Controller</li><li>• Storwize V7000</li><li>• Storwize V7000 Unified</li></ul>                   | Direct Connection                      | 443 and 22 |
| <ul style="list-style-type: none"><li>• The XIV system</li></ul>                                                                                                   | Direct Connection                      | 7778       |

Ensure that your network configuration is set up so that Tivoli Storage Productivity Center can send outgoing TCP/IP packets to the storage controllers. It is possible when adding the storage controllers to Tivoli Storage Productivity Center to set a specific port number for your storage controller.

Because there are typically multiple applications running on the management server, it is possible that port conflicts might arise if other applications attempt to use the same ports that IBM Tivoli Storage Productivity Center for Replication is using. Use the **netstat** command to verify which ports the various applications on the management server are using.

When you add a storage system to the Tivoli Storage Productivity Center for Replication configuration, the port field is automatically populated with the appropriate value. If you want to use different ports, you can change them by clicking **Storage Systems** located in the navigation tree, clicking the storage system that you want to change, and then changing the port value in the View/Modify Details panel.

**Note:** The storage system must not be in a Connected state if you want to change port values.

If firewalls are being used in your configuration, ensure that none of these ports are being blocked. Also ensure that not only is the Tivoli Storage Productivity Center for Replication server granted access to reach the other components, but that the other components are granted access to reach the Tivoli Storage Productivity Center for Replication server.

**Note:** If the management server is running Windows Server 2008, configure the inbound and outbound rules for Tivoli Storage Productivity Center. To create a new outbound rule, launch the New Outbound Rule Wizard from the Windows Firewall with Advanced Security menu.

## Storage systems

A *storage system* is a hardware device that contains data storage. Tivoli Storage Productivity Center for Replication can control data replication within and between various storage systems.

To replicate data among storage systems using Tivoli Storage Productivity Center for Replication, you must manually add a connection to each storage system in the Tivoli Storage Productivity Center for Replication configuration. This allows you to omit storage systems for which Tivoli Storage Productivity Center for Replication is not to manage replication and omit storage systems that are being managed by another Tivoli Storage Productivity Center for Replication management server.

For redundancy, you can connect a single storage system using a combination of direct, Hardware Management Console (HMC), and z/OS connections.

You can use the following storage systems:

- IBM TotalStorage Enterprise Storage Server (ESS) Model 800
- IBM System Storage DS6000
- IBM System Storage DS8000
- IBM System Storage SAN Volume Controller
- IBM Storwize V7000
- IBM Storwize V7000 Unified
- IBM XIV Storage System

A SAN Volume Controller can virtualize a variety of storage systems. Although Tivoli Storage Productivity Center for Replication does not support all storage systems, you can manage these storage systems through a single SAN Volume Controller cluster interface. Tivoli Storage Productivity Center for Replication connects directly to the SAN Volume Controller clusters.

You can define a location for each storage system and for each site in a session. When adding copy sets to the session, only the storage systems whose location matches the location of the site are allowed for selection. This ensures that a session relationship is not established in the wrong direction.

### Notes:

- Tivoli Storage Productivity Center for Replication does not automatically discover the physical locations of storage systems. You can manually assign a location to a storage system from the GUI and CLI.
- Throughout this document, *ESS/DS* refers to the following models:
  - IBM TotalStorage Enterprise Storage Server Model 800
  - IBM System Storage DS8000

## Storage connections

You must create a connection from the IBM Tivoli Storage Productivity Center for Replication management server to each storage system. You can connect either directly or through a Hardware Management Console (HMC) or IBM z/OS connection.

A single storage system can be connected using multiple connections for redundancy. For example, you can connect a IBM System Storage DS8000 storage system using an HMC connection and a z/OS connection. Tivoli Storage Productivity Center for Replication monitors how a storage system has been added to the configuration.

When you add a storage connection to the Tivoli Storage Productivity Center for Replication configuration, the storage system and the connection are added to the active management server configuration. For direct and HMC connections, the storage system and connection are also added to the standby management server configuration. For z/OS connections, only the storage system is added to the standby management server configuration. The connection is not added because the standby management server might not be running on z/OS and might not have access to the volumes on the storage system through a z/OS connection.

The storage systems are not required to be connected to the standby management server. However, if a storage system does not have a connection on the standby management server, you cannot manage copy services on the storage system from the standby server.

**Important:** If the Metro Mirror heartbeat is enabled, do not connect to a IBM TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, or System Storage DS6000 storage system using both an HMC connection and a direct connection. If you have both types of connections and the direct connection is lost, the session changes to the suspended state even though the HMC connection is still valid. If both connections are lost and the session is in the suspended state, restart the session when connectivity is regained to synchronize the session with the hardware.

When Tivoli Storage Productivity Center for Replication is running on z/OS and a storage system is added to the Tivoli Storage Productivity Center for Replication configuration through a TCP/IP (direct or HMC) connection, all ECKD volumes that are attached to the management server are managed through the TCP/IP connection. To use the Fibre Channel connection, you must explicitly add the storage system to the Tivoli Storage Productivity Center for Replication configuration through a z/OS connection.

If a storage system was previously added to the Tivoli Storage Productivity Center for Replication configuration through a z/OS connection and later the storage system is added through a TCP/IP connection, all non-attached ECKD volumes and fixed block volumes are added to the Tivoli Storage Productivity Center for Replication configuration.

When you remove a storage system, Tivoli Storage Productivity Center for Replication automatically removes all connections that the storage system is using with exception of the z/OS connection. You can also individually remove each connection through which the storage system is connected.

If Tivoli Storage Productivity Center for Replication has multiple connections to a specific storage system, the order in which you remove the connections produces different results:

- If you remove all direct and HMC connections first, the fixed block and non-attached ECKD volumes are removed from the Tivoli Storage Productivity Center for Replication configuration. The remaining ECKD volumes that are attached through the z/OS connection remain in the Tivoli Storage Productivity Center for Replication configuration until the z/OS connection is removed. Removing the TCP/IP connection also disables the Metro Mirror heartbeat.
- If you remove the z/OS connection first and if there is an HMC or direct connection to volumes, those volumes are not removed from the Tivoli Storage Productivity Center for Replication configuration.
- HyperSwap can run provided that volumes are attached and available to z/OS storage, even if you are using a TCP/IP connection to storage.

### Direct connection

The Tivoli Storage Productivity Center for Replication management server can connect directly to TotalStorage Enterprise Storage Server Model 800, DS6000, DS8000, SAN Volume Controller, Storwize V7000, Storwize V7000 Unified or the XIV system storage systems through a TCP/IP connection. The TCP/IP connection is required to discover the storage systems configuration (such as LSSs, volumes, volume size, and format), issue queries, and receive asynchronous events.

DS8000 storage systems on an IPV4 network can be connected directly to the management server. A direct connection requires an Ethernet card in the cluster. DS8000 storage systems on an IPV6 network cannot use a direct connection. They can be connected only through an HMC or z/OS connection.

When you add a direct connection to a DS or ESS cluster, specify the following information for cluster 0 and 1:

- IP addresses or domain names
- Ports
- User names
- Passwords

SAN Volume Controller or Storwize V7000 can virtualize various storage systems. Although Tivoli Storage Productivity Center for Replication does not support all storage systems, you can manage these storage systems through a single SAN Volume Controller or Storwize V7000 cluster interface. Tivoli Storage Productivity Center for Replication connects directly to the SAN Volume Controller or Storwize V7000 clusters. When you add a direct connection to a SAN Volume Controller or Storwize V7000 cluster to the Tivoli Storage Productivity Center for Replication configuration, specify the cluster IP address of the SAN Volume Controller or Storwize V7000 cluster, which in turn points to multiple SAN Volume Controller or Storwize V7000 storage systems. Ensure that the user name and password are correct for the cluster. If incorrect values are used, significant communication problems can occur, such as never advancing to the Prepared state.

**Important:** The SAN Volume Controller or Storwize V7000 user name must have privileges to maintain SSH keys. For information about troubleshooting Secure Shell connections to the SAN Volume Controller or Storwize V7000, see the Ethernet Connection Restrictions on SAN Volume Controller website at [www-01.ibm.com/support/docview.wss?uid=ssg1S1002896](http://www-01.ibm.com/support/docview.wss?uid=ssg1S1002896).

## Hardware Management Console connection

The IBM Tivoli Storage Productivity Center for Replication management server can connect to DS8000 storage systems through a Hardware Management Console (HMC). An HMC can have multiple DS8000 storage systems connected to it. When you add an HMC to the IBM Tivoli Storage Productivity Center for Replication configuration, all DS8000 storage systems that are behind the HMC are also added. You cannot add or remove individual storage systems that are behind an HMC.

You can also add a dual-HMC configuration, in which you have two HMCs for redundancy. This is recommended when the Metro Mirror heartbeat is required. You must configure both HMCs identically, including the user ID and password.

If planned maintenance is necessary on the HMC, it is recommended that you disable the Metro Mirror heartbeat on the management server while the maintenance is performed.

If the HMC needs to go down frequently or reboots frequently, it is recommended that you disable the Metro Mirror heartbeat. If the Metro Mirror heartbeat is required, the direct connection is recommended instead of an HMC connection.

**Important:** If a DS8000 storage system uses an HMC connection, the Metro Mirror heartbeat could trigger a freeze on the storage system and impact applications for the duration of the long busy timeout timer if the HMC is shut down for any reason, including upgrading microcode. The long busy timeout timer is the time after which the storage system will allow I/O to begin again after a freeze occurs if no run command has been issued by IBM Tivoli Storage Productivity Center for Replication. The default value is two minutes for ECKD volumes or one minute for fixed block volumes.

### Notes:

- The user ID that you use to connect to the HMC must have admin, op\_storage, or op\_copy\_services privileges on the DS8000 storage system.
- For minimum microcode requirements to connect to a DS8000 through a management console, see the Supported Storage Products List website at [www-01.ibm.com/support/docview.wss?uid=swg21386446](http://www-01.ibm.com/support/docview.wss?uid=swg21386446).

## z/OS connections

An IBM Tivoli Storage Productivity Center for Replication management server that runs on z/OS can connect to IBM TotalStorage Enterprise Storage Server (ESS) Model 800, DS8000, and DS6000 storage systems through a z/OS connection. The z/OS connection is used to issue replication commands and queries for attached ECKD volumes over an existing Fibre Channel network and to receive asynchronous events. When a storage system is added to IBM Tivoli Storage Productivity Center for Replication through the z/OS connection, all ECKD volumes that are attached to the IBM Tivoli Storage Productivity Center for Replication management system are added to the IBM Tivoli Storage Productivity Center for Replication configuration. ECKD volumes that are not attached to the IBM Tivoli Storage Productivity Center for Replication z/OS management server are not added to the IBM Tivoli Storage Productivity Center for Replication configuration through the z/OS connection.

### Notes:

- Ensure that all volumes in the logical storage subsystem (LSS) that you want to manage through a z/OS connection are attached to z/OS. Either the entire LSS



must be attached to z/OS or none of the volumes in the LSS should be attached to z/OS for IBM Tivoli Storage Productivity Center for Replication to properly manage queries to the hardware.

- The z/OS connection is limited to storage systems that are connected to an IBM Tivoli Storage Productivity Center for Replication management server running z/OS.
- The Metro Mirror heartbeat is not supported through the z/OS connection. To use the Metro Mirror heartbeat, the storage systems must be added using a direct connection or Hardware Management Console (HMC) connection. If the Metro Mirror heartbeat is enabled, a storage system is added through a direct connection and z/OS connection, and the direct connection becomes disconnected, then a suspend results as there is no heartbeat through the z/OS connection.

If at least one volume in a Logical Storage Subsystem (LSS) is attached through a z/OS connection, then all volumes in that LSS must be similarly attached. For example, if there are two ECKD volumes in an LSS, and one volume is attached to the IBM Tivoli Storage Productivity Center for Replication system using a z/OS connection and the other is attached through a direct connection, IBM Tivoli Storage Productivity Center for Replication would have knowledge of direct-connected volume. IBM Tivoli Storage Productivity Center for Replication issues commands to both volumes over the Fibre Channel network; however, commands issued to the direct-connection volume will fail, and IBM Tivoli Storage Productivity Center for Replication will show that the copy set that contains that volume has an error.

Use the following guidelines to add storage systems through a z/OS connection:

- Use the z/OS connection to manage ECKD volumes that are attached to an IBM Tivoli Storage Productivity Center for Replication management server running z/OS.
- To manage z/OS attached volumes through a z/OS connection (for example, for HyperSwap), you must explicitly add the z/OS connection for that storage system in addition to a TCP/IP connection (either the direct connection or the HMC connection).
- Create a z/OS connection before all TCP/IP connections if you want to continue to have IBM Tivoli Storage Productivity Center for Replication manage only the attached ECKD volumes.

**Tip:** It is recommended that you create both TCP/IP and z/OS connections for ECKD volumes to allow for greater storage accessibility.

## Protected volumes

You can mark volumes as protected if you do not want those volumes used for replication.

When a volume is marked as protected, you cannot include that volume in a copy set. This protection applies only to IBM Tivoli Storage Productivity Center for Replication.

You might want to protect a volume in the following instances:

- The volume contains data that you never want to be copied to another volume. For example, the volume is secure, but if the data is copied to an unsecured volume, the data could be read. For this reason, the volume should not be the source for a relationship.

- The volume contains data that you do not want to be overwritten. For this reason, the volume should not be the target of a relationship.

Only administrators can change the volume protection settings.

## Site awareness

You can associate a location with each storage system and each site in a session. This *site awareness* ensures that only the volumes whose location matches the location of the site are allowed for selection when you add copy sets to the session. This prevents a session relationship from being established in the wrong direction.

**Note:** To filter the locations for site awareness, you must first assign a site location to each storage system.

IBM Tivoli Storage Productivity Center for Replication does not perform automatic discovery of locations. Locations are user-defined and specified manually.

You can change the location associated with a storage system that has been added to the IBM Tivoli Storage Productivity Center for Replication configuration. You can choose an existing location or add a new one. Locations are deleted when there is no longer a storage system with an association to that location.

When adding a copy set to a session, a list of candidate storage systems is presented, organized by location. Storage systems that do not have a location are displayed and available for use when you create a copy set.

You can also change the location for any site in a session. Changing the location of a session does not affect the location of the storage systems that are in the session.

Changing the location of a storage system might have consequences. When a session has a volume role with a location that is linked to the location of the storage system, changing the location of the storage system could change the session's volume role location. For example, if there is one storage system with the location of A\_Location and a session with the location of A\_Location for its H1 role, changing the location of the storage system to a different location, such as B\_Location, also changes the session's H1 location to Site 1. However, if there is a second storage system that has the location of A\_Location, the session's role location is not changed.

**Important:** Location matching is enabled only when adding copy sets. If you change the location of a storage system or volume role, IBM Tivoli Storage Productivity Center for Replication does not audit existing copy sets to confirm or deny location mismatches.

## Adding a storage connection

You can add one or more connections to a storage system. To replicate data among storage systems by using Tivoli Storage Productivity Center for Replication, you must add connections to the storage systems.

**Prerequisites:** You must have Administrator privileges to add a storage connection.

A single storage system can be connected by a combination of direct, management-console, and z/OS connections. Perform these steps to add a storage system connection:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Click **Add Storage Connection**. The Add Storage System wizard Type page is displayed.
3. Select the connection type that you want to use to add the storage systems and click **Next**. The Connection page is displayed.
4. Perform one of these steps depending on the type of storage connection that you selected:

- **DS8000 / ESS 800 / DS6000 (Direct Connection)**: Enter the IP address or domain name, port number, user name, and password for both cluster 0 and cluster 1 in the appropriate fields.

In a System Storage DS8000 environment, if resource groups are defined on the System Storage DS8000 that you are connecting to, the user name you define here must have the appropriate access level on the System Storage DS8000 to manage copy services for the volumes that are used by Tivoli Storage Productivity Center for Replication.

- **DS8000 (HMC Connection)**: Enter the primary HMC IP address or domain, user name, password, and optionally the secondary HMC IP address or domain name in the appropriate fields.
- **SAN Volume Controller / Storwize V7000 / Storwize V7000 Unified (Direct Connection)**: Enter the cluster IP address or domain name, user name, and password in the appropriate fields.

The following are considerations for connection:

- a. Authentication with the storage system cluster is performed by public key exchange. When adding a storage system cluster to the Tivoli Storage Productivity Center for Replication configuration, specify the user ID and password of an administrator on the cluster that has sufficient privileges to maintain the Secure Shell (SSH) keys.
  - b. When a valid user name and password are specified, Tivoli Storage Productivity Center for Replication audits the public keys installed on the storage system cluster to ensure that the Tivoli Storage Productivity Center for Replication key has the correct access levels. If the key is not installed on the cluster, Tivoli Storage Productivity Center for Replication attempts to install the key. If you install the key yourself, install the public key found in the *install\_root/eWAS/profiles/CSM/etc/public* directory on the storage system cluster by using the SSH Key Maintenance panel on the storage system graphical user interface. The public key ID is installed under tpcr.
  - c. You must use a storage system cluster superuser user name and password to correctly receive indications from the storage system cluster.
  - d. For Storwize V7000, Storwize V7000 Unified, or SAN Volume Controller version 5.0 or later, you must create a user ID through the storage system administration console on the cluster to which you want to attach Tivoli Storage Productivity Center for Replication. This user ID must have local authentication type and have the Administrator role authority.
- **XIV (Direct Connection)**: Enter the IP address or domain name of the XIV system node. Enter a user name and password of a user with appropriate access rights to the XIV system.
  - **z/OS (FICON Connection)**: Select one or more storage systems from the list of candidate systems that match the selected storage type, or click **All** to select all storage systems.
5. Click **Next**.

6. Click **Finish**.

## Removing a storage connection

You can remove a single connection to a storage system from the IBM Tivoli Storage Productivity Center for Replication configuration.

**Prerequisites:** You must have Administrator privileges to remove a storage connection.

When removing a connection, the storage system might have other connections and therefore still be connected to IBM Tivoli Storage Productivity Center for Replication.

When removing a connection, all storage systems that rely on only that connection are removed. If a storage system is removed, the volumes for that storage system are removed from management server control. All copy sets with a volume on the removed storage systems are also removed from their respective sessions, making the target volume unrecoverable.

If Tivoli Storage Productivity Center for Replication has multiple connections to a specific storage system, the order in which you remove the connections produces different results:

- If you remove all direct and HMC connections first, the fixed block and non-attached ECKD volumes are removed from the Tivoli Storage Productivity Center for Replication configuration. The remaining ECKD volumes that are attached through the z/OS connection remain in the Tivoli Storage Productivity Center for Replication configuration until the z/OS connection is removed. Removing the TCP/IP connection also disables the Metro Mirror heartbeat.
- If you remove the z/OS connection first and if there is an HMC or direct connection to volumes, those volumes are not removed from the Tivoli Storage Productivity Center for Replication configuration.
- HyperSwap can run provided that volumes are attached and available to z/OS storage, even if you are using a TCP/IP connection to storage.

Perform one of these procedures to remove a storage connection:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Click the **Connections** tab.
3. Select the storage connection that you want to remove.

**Important:** If you choose to delete an HMC connection, all storage systems that share the HMC connection will also be removed.

4. Select **Remove Connection** from the **Actions** list, and click **Go**.
5. Click **Yes** to remove the storage system.

## Removing a storage system

You can remove a storage system from the IBM Tivoli Storage Productivity Center for Replication configuration.

**Prerequisites:** You must have Administrator privileges to remove a storage system.

Removing a storage system removes all volumes on that storage system from management server control. All copy sets with a volume on the removed storage

system are removed from their respective sessions, making the target volume unrecoverable. All connections to the removed storage system are removed, and any storage systems sharing these connections are also removed.

Perform these steps to remove a storage system:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Select the storage system that you want to remove.

**Important:** All connections to this storage system will be removed, all volumes on the storage system will be removed from management server control, and all copy sets that have a volume on this storage system will be removed from their respective sessions, leaving the target volume unrecoverable. Any storage systems sharing these connections will be removed as well.

3. Select **Remove Storage System** from the **Actions** list, and click **Go**.
4. Click **Yes** to remove the storage system.

## Modifying the location of storage systems

You can associate a location with a storage system after a connection has been made to that storage system.

**Prerequisites:** You must have Administrator privileges to modify the location of a storage system.

Changing the location of a storage system might have consequences. When a session has a volume role with a location that is linked to the location of the storage system, changing the location of the storage system could change the session's volume role location. For example, if there is one storage system with the location of A\_Location and a session with the location of A\_Location for its H1 role, changing the location of the storage system to a different location, such as B\_Location, also changes the session's H1 location to Site 1. However, if there is a second storage system that has the location of A\_Location, the session's role location is not changed.

Perform these steps to modify the location of a storage system:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Change the location of the storage system by selecting a previously defined location from the drop-down list or type a new name in the table cell.

To disable site awareness, set the location to None.

**Note:** Locations are deleted from the drop-down list when there is no longer a storage system with an association to that location.

## Modifying storage connection properties

You can modify the connection properties for a storage system, including IP addresses, user name and password.

**Prerequisites:**

- You must have Administrator privileges to modify storage connection properties.
- The storage system must be in the Disconnected state to change most storage connection parameters. You can add a secondary HMC to an existing HMC connection without the HMC being disconnected.

A storage system can lose connection to the management server, for example, if a port is blocked by a firewall or the user name or password is changed on the storage system. If the storage system loses connection, you might need to modify parameters (for example, user name or password) manually on the storage system, and then update the parameters in IBM Tivoli Storage Productivity Center for Replication.

Perform these steps to modify storage connection properties:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Click the **Connections** tab.
3. Perform one of these steps to view details for a specific storage connection:
  - Click the storage connection ID.
  - Select the storage connection, click **View/modify Connection Details** from the actions list, and then click **Go**.
4. Modify the appropriate settings match the settings for the storage system.
5. Click **Apply** to continue making changes, and click **OK** when finished.

## Refreshing the storage system configuration

You can refresh the storage system configuration to query the storage system for changes, such as which volumes are contained in an LSS. You might do this when you reconfigure a storage system and you want IBM Tivoli Storage Productivity Center for Replication to be aware of the changes.

**Prerequisites:** You must have Administrator privileges to modify storage connection settings.

Perform these steps to refresh the storage configuration:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Select the storage system for which you want to refresh the configuration.
3. Select **Refresh Configuration** from the **Actions** list, and click **Go**.

## Setting volume protection

To ensure that data on a volume is not overwritten, you set its status to protected. Protected volumes are excluded from replication.

You must have Administrator privileges to change the protection setting of a volume.

1. In the navigation tree, select **Storage Systems**. The Storage Systems page is opened.
2. Click **Volume Protection**. The Volume Protection wizard is opened.
3. Select a storage system.
4. Optional: Depending on the type of storage system:
  - a. Select **All IO Groups** or a specific I/O group.
  - b. Select **All Logical Storage Subsystems** or a specific logical storage subsystem.
  - c. Select **All Pools** or a specific pool.
5. Optional: In the **Volume** field, select a single volume.

6. Optional: In the **Volume Mask** field, enter a sequence of characters and wildcards that match user-defined or system-defined volume IDs. To protect a specific volume, enter the volume ID such as ESS:2105.FCA57:VOL:1000. To use a pattern to retrieve one or more volume IDs, you can enter a partial volume ID and use the wildcard character (\*) to represent zero or more characters. For example, to retrieve all volume IDs that contain the characters FCA57, you enter \*FCA57\*.
7. Click **Next**.
8. Verify the search results, and click **Next**.
9. Click **Select All** to protect all the volumes. Alternatively, select a check box next to the volumes that you want to protect.
10. Click **Next**.
11. Click **Finish**.

## Restoring data from a journal volume

This topic provides information on restoring data from a journal (J) volume. It describes how to restore data from journal volume used as part of a ESS/DS6000/DS8000 Global Mirror session, or as part of a Metro Global Mirror session, if data was corrupted on a host volume after you issued a **Recover** command. Following these steps will enable you to return to a consistent copy of the data on the host volume.

Perform the following steps to move the data from the journal volume back to the host volume:

**Note:** Follow these instructions only if you have already issued a **Recover** command to the site containing the journal volume. After the **Recover** command is issued, the journal volume will hold a copy of the consistent data at the time the command was issued.

1. Outside of IBM Tivoli Storage Productivity Center for Replication, using the DS8000 GUI /CLI, issue withdraw initiate background copy (issue a **rmflash -cp** command) on pairs containing the journal volume (for example, H2J2). This copies the remaining uncopied tracks from the host to the journal. Then, ensure all the Out of Sync (OOS) tracks reach zero.
2. Create a separate FlashCopy session either with IBM Tivoli Storage Productivity Center for Replication, or with the DS8000 GUI /CLI (issue a **mkflash** command with background copy), with the following conditions:
  - The journal volume (Jx) is the source volume.
  - The host volume (Hx or Ix if using a session with Practice capabilities) is the target volume.
  - x is the site the **Recover** command was issued to.

---

## Managing host systems

*Host system* refers to an IBM AIX 5.3 or 6.1 server that is connected to IBM System Storage DS8000 devices. A connection from Tivoli Storage Productivity Center for Replication to the host system is used to enable the automatic swap of input/output (I/O) from the primary storage unit to the secondary storage unit in the case of a primary error.

A connection to the host system is required to use the Tivoli Storage Productivity Center for Replication Open HyperSwap feature. For the software and hardware required to support Open HyperSwap, see “Setting up the environment for Open HyperSwap” on page 101.

Connecting to a host system requires the IP address or host name of the system and the port number for communication. All connections are secured via a Secure Socket Layer (SSL) connection.

**Restriction:** Open HyperSwap is not supported for AIX host servers that are in a clustered environment such as PowerHA (previously known as HACMP).

**Related tasks:**

“Managing a session with HyperSwap and Open HyperSwap replication” on page 100

HyperSwap and Open HyperSwap provide high availability of data in the case of a primary disk storage system failure. When a failure occurs in writing input/output (I/O) to the primary storage system, the failure is detected by IOS, and IOS automatically swaps the I/O to the secondary site with no user interaction and little or no application impact.

## Adding a host system connection

You can add a connection to one or more host systems to the IBM Tivoli Storage Productivity Center for Replication configuration.

**Prerequisites:** You must have Administrator privileges to add a host system connection.

For the software and hardware required to support Open HyperSwap, see “Setting up the environment for Open HyperSwap” on page 101.

Perform these steps to add a host system connection:

1. In the navigation tree, select **Host Systems**. The Host Systems panel is displayed.
2. Click **Add Host Connection**. The Add Host Connection dialog box is displayed.
3. Enter the host name or IP address and the port for the host system and click **Add Host**. The host system is displayed on the Host Systems panel. The default port is 9930. Unless the port has been modified in Subsystem Device Driver Path Control Module (SDDPCM), use the default port.

The host system is displayed in the Host System table. If the connection is successful, the status **Connected** is displayed for the connection. If the connection was not successful, the status **Disconnected** is displayed.

## Modifying a host system connection

You can modify host system connections in the IBM Tivoli Storage Productivity Center for Replication configuration.

**Prerequisites:** You must have Administrator privileges to modify a host system connection and the connection must be in a disconnected state.

Perform these steps to modify a host system connection:



1. In the navigation tree, select **Host Systems**. The Host Systems panel is displayed.
2. Select the host system connection that you want to modify.
3. Select **Modify Host Connection** from the **Select Action** list, and click **Go**.
4. Modify the information that is presented for the host system and click **Update Host**. The updated host system information is displayed on the Host Systems panel.

The updated host system information is displayed in the Host System table. If the connection is successful, the status **Connected** is displayed for the connection. If the connection was not successful, the status **Disconnected** is displayed.

## Removing a host system connection

You can remove host system connections from the IBM Tivoli Storage Productivity Center for Replication configuration.

**Prerequisites:** You must have Administrator privileges to remove a host system.

Removing a host system connection disables the ability to use Open HyperSwap. Any session using the host system to provide Open HyperSwap capabilities can no longer communicate with the host and Open HyperSwap is disabled for the entire session.

Perform these steps to remove a host system connection:

1. In the navigation tree, select **Host Systems**. The Host Systems panel is displayed.
2. Select the host system connection that you want to remove.
3. Select **Remove Host Connection** from the **Select Action** list, and click **Go**.
4. Click **OK** to remove the host system connection.

## Removing a session from a host system connection

You can remove a session that is associated with a host system from the IBM Tivoli Storage Productivity Center for Replication configuration. Once removed, the host no longer recognizes the session that is managing the volumes attached to that host. This function is to be used primarily for cleanup purposes.

When a session has Open HyperSwap enabled, the session communicates with the host system and the host system stores an association to that session on the IBM Tivoli Storage Productivity Center for Replication server. If the Tivoli Storage Productivity Center for Replication server that made the association becomes inaccessible, it might be necessary to clean up and remove the session association from a different Tivoli Storage Productivity Center for Replication server.

If a host system has an associated session, the session name is displayed in the **Sessions** column of the Host Systems table. If the session is a session that is currently defined on the Tivoli Storage Productivity Center for Replication server, the session name is displayed as a link. The link opens the Session Details panel.

If the session name is not a session on the server, an icon is displayed. This session must be removed because the host system can support only a single session association. The session association must be removed before a Tivoli Storage Productivity Center for Replication server can re-establish capabilities with the host system.

**Prerequisites:** You must have Administrator privileges to remove a host system.

Perform these steps to remove session association from a host system connection:

1. In the navigation tree, select **Host Systems**. The Host Systems panel is displayed.
2. Select the host system connection that contains the session that you want to remove.
3. Select **Remove Session Association** from the **Select Action** list, and click **Go**.
4. Click **OK** to remove the host system connection.

The session name is removed from the **Session** column on the Host Systems table.

---

## Managing logical paths

Logical paths define the relationship between a source logical subsystem (LSS) and a target LSS that is created over a physical path. To configure logical paths for TotalStorage Enterprise Storage Server, System Storage DS8000, and DS6000, use the ESS/DS Paths panel in Tivoli Storage Productivity Center.

To configure partnerships for the following storage systems, use the graphical user interface (GUI) or command-line interface (CLI):

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified
- XIV system

## Viewing logical paths

You can view all logical paths that are defined on an IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, or IBM System Storage DS6000 storage system.

Perform one of these procedures to view logical paths:

- From the ESS/DS Paths panel of IBM Tivoli Storage Productivity Center for Replication:
  1. In the navigation tree, select **ESS/DS Paths**. The ESS/DS Paths panel is displayed.
  2. Click the storage system ID to display logical paths for that storage system.
- From the Storage Systems panel:
  1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
  2. Select an ESS, DS6000, or DS8000 storage system for which you want to view logical paths.
  3. Select **View Paths** from the **Select Action** list, and click **Go**. The ESS/DS Paths panel is displayed with a list of defined logical paths.

## Adding logical paths

This topic describes how to add IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, and IBM System Storage DS6000 logical paths.

Ensure that you have defined the appropriate storage systems on the Storage Systems panel.

Perform these steps to add logical paths:

1. In the IBM Tivoli Storage Productivity Center for Replication navigation tree, select **ESS/DS Paths**.
2. Click **Manage Paths**. The Path Management wizard is displayed.
3. From the drop-down boxes in the Path Management wizard, select the source storage system, source logical storage system, target storage system, and target logical storage system. Then, click **Next**.
4. From the drop-down boxes in the Path Management wizard, select the source port and target port and click **Add**. You can add multiple paths between the logical storage subsystems, or just one at a time. When you have made your selections, click **Next**.
5. Confirm your selections and click **Next**.
6. Verify the remaining wizard panels and click **Next**.
7. Click **Finish**.

## Adding logical paths using a CSV file

You can create a comma separated (CSV) file to define logical paths. The CSV file specifies storage systems pairings and associated port pairings that are used for replication. IBM Tivoli Storage Productivity Center for Replication uses the port pairings defined in the CSV file to establish logical paths

Perform these steps to add IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, and IBM System Storage DS6000 logical paths using a CSV file:

1. Create a CSV file named `portpairings.csv` in the `websphere_home/AppServer/profiles/websphere_profile/properties` directory. You can create the CSV file in a spreadsheet such as Microsoft Excel or in a text editor. An example CSV file is as follows:

```

Example CSV file

2107.04131:2107.01532,0x0331:0x0024,0x0330:0x0100,0x0331:0x000C
2107.05131:2107.01532,0x0330:0x0029,0x0331:0x0001
```

Each line represents a storage system to storage system pairing. The first value represents the storage systems, which are delimited by a colon. The remaining values are the port pairs, which are delimited by a colon. All values are separated by a comma. Commented lines must start with #,

2. To enable the changes in the file, perform a task that requires new paths to be established. For example, suspend a session to remove the logical paths and then issue the **Start H1->H2** command to enable the paths to use the port pairings in the CSV file.

### Considerations when creating and using the CSV file:

- The CSV file does not affect Global Mirror control paths.
- Port mapping is bi-directional. A logical path is established from system A to system B and from system B to system A depending on the direction of the pairs on the hardware.
- If the CSV file contains multiple lines that specify the same storage system to storage system pairing, Tivoli Storage Productivity Center for Replication uses the last line. This rule applies regardless of the order of the storage system pairing. For example, if you have storage systems 2107.04131:2107.01532 defined

on the first line of the CSV file and then have 2107.01532:2107.04131 defined on the second line, Tivoli Storage Productivity Center for Replication uses second line.

- If a line in the CSV contains information that is not formatted correctly, the line is ignored. This rule includes lines that specify storage systems but do not include ports or include ports that are not formatted correctly.
- If the CSV file contains valid and invalid port pairs, the valid port pairs might or might not be established. Invalid port pairs can cause the following errors to be displayed in the Tivoli Storage Productivity Center for Replication console and in the ESS/DS Paths panel:
  - Return Code F52: This error is displayed if a port is invalid.
  - Return Code 0400: This error is displayed if a port is invalid and out of the range for the device.

Other storage system error codes might be displayed also, depending on the path topology, types of paths, and the incorrect port pairings specified in the CSV file.

- If the CSV file contains no valid port pairs, no logical paths are established and subsequent commands to the storage systems that require logical paths might fail. If there are existing logical paths for a storage system, those paths are used until they are removed.

## Removing logical paths

This topic describes how to remove IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, and IBM System Storage DS6000 logical paths.

1. In the IBM Tivoli Storage Productivity Center for Replication navigation tree, select **ESS/DS Paths**.
2. Click the link for the storage system that contains the paths that you want to remove.
3. Select the paths that you want to remove.
4. From the drop-down box, select **Remove**.
5. Click **Go**.

---

## Setting up data replication

This topic describes the how to set up data replication in your environment, including creating sessions and adding copy sets to those sessions.

A *session* is a container of multiple copy sets managed by a replication manager. A *copy set* is a set of volumes that contain copies of the same data. All the volumes in a copy set are the same format (count key data [CKD] or fixed block) and size. In a replication session, the number of volumes in a copy set and the role that each volume in the copy set plays are determined by the session type.

## Sessions

A *session* is used to perform a specific type of data replication against a specific set of volumes. The source volume and target volumes that contain copies of the same data are collectively referred to as a copy set. A session can contain one or more copy sets.

If a session has failover and failback capabilities, you can perform a site switch in which you move the application from one site to another and change the direction of the copy without having to perform a full copy. Without failover and failback

capabilities, each time you move the application and writing to a different site in the session, you must initiate a full copy to synchronize the new source with the new target to regain disaster recovery capability. An IBM Tivoli Storage Productivity Center for Replication session with failover and failback capabilities uses the hardware's ability to track changes after a suspension (where applicable), so only the changed data must be resynchronized.

The type of data replication (also known as the session type) that is associated with the session determines the actions that can be performed against all copy sets in the session, the number of volumes in each copy set, and the role that each volume plays.

**Important:** Use only the Tivoli Storage Productivity Center for Replication GUI or CLI to manage session relationships, such as volume pairs and copy sets. Do not modify session relationships through individual hardware interfaces, such as the DSCLI. Modifying relationships through the individual hardware interfaces can result in a loss of consistency across the relationships managed by the session, and might cause the session to be unaware of the state or consistency of the relationships.

### Copy sets

During data replication, data is copied from a source volume to one or more target volumes, depending on the session type. The source volume and target volumes that contain copies of the same data are collectively referred to as a *copy set*.

Each volume in a copy set must be of the same size and machine type (for example, 3380 volumes must be used with other 3380 volumes and SAN Volume Controller volumes must be used with other SAN Volume Controller volumes). The number of volumes in the copy set and the role that each volume plays is determined by the session type (or copy type) that is associated with the session to which the copy set belongs.

Use the Add Copy Sets wizard to add copy sets to an existing session. You can select a storage system; a logical subsystem (LSS), I/O group, or pool; or single volume for each role and then create one or more copy sets for the session.

You can use one of the following volume pairing options to automatically create multiple copy sets in the same session.

#### **Storage system matching (System Storage DS8000, System Storage DS6000, or TotalStorage Enterprise Storage Server Model 800 Metro Mirror sessions only)**

Creates copy sets by matching volumes (based on the volume IDs) across all logical subsystems (LSSs) for the selected storage systems. For example, volume 01 on the source LSS is matched with volume 01 on the target LSS.

You cannot select the storage system and select **All Logical Subsystems** in the list of LSSs. You can also do auto-matching at the LSS level for Metro Mirror sessions.

#### **LSS, I/O group, or pool matching**

Creates copy sets by matching all volumes based on the selected LSS, I/O group, or pool for each role in copy set.

Select the storage system and LSS, I/O group, or pool, and then select **All Volumes** in the **Volume** list.

If you do not want to use the auto-generated volume pairing for a copy set, clear that copy set so that it is not added during the wizard. Then, add the remaining copy sets and reopen the Add Copy Set wizard and manually enter the volume pairings that you want.

Invalid copy set are not added to the session. Copy sets can be invalid if their volumes are not the same type or size.

You can remove copy sets that you do not want to add to the session, even if they valid. This enables you to filter and eliminate unwanted copy sets before they are added to the session.

You can export the copy sets to take a snapshot of your session at a particular point in time for backup purposes.

**Note:** You can copy an entire storage system only for Metro Mirror sessions.

### **Considerations for adding copy sets**

When you create a copy set, a warning is displayed if one or more selected volumes already exist in another session. Whether is safe to add the created copy set to the session depends on the environment. For example, if you created one session for normal replication and another session for a disaster recovery practice scenario, you must use the same target volumes from the original session as the source volumes in the practice session. If the volume you selected is already in another session, confirm that this is the configuration you want.

You can use extent space-efficient volumes as copy set volumes for Global Mirror with Practice sessions for System Storage DS8000 6.3 or later. Extent space-efficient volumes must be fixed block (FB). You cannot use count key data (CKD) volumes.

You can use extent space-efficient volumes as source, target, and journal volumes. If you use an extent space-efficient volume as a source or target volume in the copy set, you must use extent space-efficient volumes for all source and target volumes in the copy set. In this situation, the journal volumes can be extent space-efficient volumes, track space-efficient volumes, or a combination of both volume types. If extent space-efficient volumes are not used as source or target volumes, journal volumes can be extent space-efficient, track space-efficient, and other types of volumes.

### **Considerations for removing copy sets**

You remove a copy set or range of copy sets by selecting the source volume; LSS, I/O group, or pool; or storage system. When the list of copy sets that meet your criteria is displayed, you can select the copy sets that you want to remove.

The consequence of removing copy sets varies depending on the state of the session:

#### **Defined**

There is no relationship on the hardware. This removes the copy set from Tivoli Storage Productivity Center for Replication data store.

#### **Preparing or Prepared**

The copy set is currently copying data, so Tivoli Storage Productivity Center for Replication terminates the hardware relationship for the copy set. The rest of the copy sets continue to run uninterrupted.

### **Suspended or Target Available**

Any existing relationships on the hardware are removed for the copy set.

Before removing all copy sets from that session, terminate the session. Removing the copy sets when the session is active can considerably increase the amount of time it takes for the copy set removal to complete. Copy sets are removed one at a time, and when the session is active, this requires commands being issued to the hardware. However, if you terminate the session first, then commands are not issued to the hardware and the removal process completes faster.

**Tip:** When you a remove copy set from Tivoli Storage Productivity Center for Replication, you might want keep hardware relationships on the storage systems. These relationships are useful when you want to migrate from one session type to another or when resolving problems. For more information about keeping the hardware relationships when removing copy sets, see Removing Copy Sets.

The behavior that occurs when a copy set is removed varies depending on the storage system:

#### **ESS 800, DS6000, and DS8000:**

- The complete copy set is removed from Tivoli Storage Productivity Center for Replication.
- Any peer-to-peer remote copy (PPRC) pair that is part of a Global Mirror consistency group is removed from the consistency group on the storage system.
- If the PPRC pair is part of a Global Mirror consistency group and is the last remaining source volume in a subordinate session, the subordinate session is removed from the storage system.
- If the PPRC pair is the last remaining participant in a Global Mirror session, the Global Mirror session is removed from the storage system.
- Any PPRC relationship remains on the storage system.
- A Metro Mirror (synchronous PPRC) pair that is in a HyperSwap configuration is removed from that configuration but the pair remains on the hardware.
- FlashCopy relationship remains on the storage system if the hardware has not already completed any background copy.

#### **SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, or the XIV system**

- The complete copy set is removed from Tivoli Storage Productivity Center for Replication.
- FlashCopy, Metro Mirror, and Global Mirror relationships are pulled out of their consistency group. If they are the last remaining relationship in a consistency group, that consistency group is removed from the hardware.

When you specify the force removal option, all knowledge of the specified copy set is removed from Tivoli Storage Productivity Center for Replication, even if the relationship itself still exists. If this occurs, you are not able to remove the relationship using Tivoli Storage Productivity Center for Replication, because no information about the relationship exists. If you force a removal of a copy set and the removal fails, you must manually remove the relationship from the hardware. If you do not, you cannot to create new relationships.

One benefit of forcing a removal of the copy sets is that Tivoli Storage Productivity Center for Replication does not manage the consistency of copy sets that it has no knowledge of. This means that additional commands to the session do not affect the removed copy sets, even though they are still in a relationship on the hardware.

If you do not specify the force removal option and an error occurs that prevents the hardware relationships from being removed, the copy set will not be removed from Tivoli Storage Productivity Center for Replication. The copy set remains as part of the session, and you can still perform actions on it.

To re-add the copy set to the session, you must perform a full copy of the data.

### **Volume roles**

Volume roles are given to every volume in the copy set. The role defines how the volume is used in the copy set and, for multi-site sessions, the site location of the volume. For example, the H1 role is made up of host-attached volumes that are located at the primary site.

The site determines the location of the volumes. The number of sites in a copy set is determined by the type of data replication (also known as the session type) that is associated with the session. IBM Tivoli Storage Productivity Center for Replication supports up to three sites:

- Site 1** The location of the primary storage that contain the source data. Upon initial configuration, this site contains the host volumes with updates that are copied to the target volumes.
- Site 2** The location of the secondary storage that receives the copy updates from the primary storage.
- Site 3** (Metro Global Mirror only) The location of the tertiary storage that receives data updates from the secondary storage.

The volume roles that are needed in a copy set are determined by the type of replication that is associated with the session. IBM Tivoli Storage Productivity Center for Replication supports these volume roles:

#### **Host volume**

A volume that is connected to a server that reads and writes I/O. A host volume can be the source of updated tracks when the server connected to the host volume is actively issuing read and write input/output (I/O). A host volume can also be the target of the replication. When the host volume is the target, writes are inhibited.

Host volumes are abbreviated as H $x$ , where  $x$  identifies the site.

#### **Journal volume**

A volume that stores data that has changed since the last consistent copy was created. This volume functions like a journal and holds the required data to reconstruct consistent data at the Global Mirror remote site. When a session must be recovered at the remote site, the journal volume is used to restore data to the last consistency point. A FlashCopy replication is created between the host or intermediate volume and the corresponding journal volume after a recover request is initiated to create another consistent version of the data.

Journal volumes are abbreviated as J $x$ , where  $x$  identifies the site.



### Intermediate volume

A volume that receives data from the primary host volume during a replication with practice session. During a practice, data on the intermediate volumes is flash copied to the practice host volumes.

Depending on the replication method being used, data on intermediate volumes might not be consistent.

Intermediate volumes are abbreviated as Ix, where x identifies the site.

### Target volume

(FlashCopy only) A volume that receives data from a source, either a host or intermediate volume. Depending on the replication type, that data might or might not be consistent. A target volume can also function as a source volume. For example, a common use of the target volume is as a source volume to allow practicing for a disaster (such as data mining at the recovery site while still maintaining disaster recovery capability at the production site).

### Role pairs

A *role pair* is the association of two volume roles in a session that take part in a copy relationship. For example, in a Metro Mirror session, the role pair can be the association between host volumes at the primary site and host volumes at the secondary site (H1-H2).

The flow of data in the role pair is shown using an arrow. For example, H1>H2 denotes that H1 is the source and H2 is the target.

*Participating role pairs* are role pairs that are currently participating in the session's copy.

*Non-participating role pairs* are role pairs that are not actively participating in the session's copy.

Snapshot sessions do not use role pairs.

### Practice volumes

You can use a *practice volume* to practice what you would do in the event of a disaster, without interrupting current data replication. Practice volumes are available in Metro Mirror, Global Mirror, and Metro Global Mirror sessions.

To use the practice volumes, the session must be in the prepared state. Issuing the Flash command against the session while in the Prepared state creates a usable practice copy of the data on the target site.

**Note:** You can test disaster-recovery actions without using practice volumes; however, without practice volumes, you cannot continue to copy data changes between volumes while testing disaster-recovery actions.

### Consistency groups

For Global Mirror and Metro Global Mirror sessions, IBM Tivoli Storage Productivity Center for Replication manages the consistency of dependant writes by creating a consistent point-in-time copy across multiple volumes or storage systems. A *consistency group* is a set of target volumes in a session that have been updated to preserve write order and are therefore recoverable.

*Data exposure* is the period when data is written to the storage at the primary site until data is replicated to storage at the secondary site. Data exposure is influenced by factors such as:

- Requested consistency-group interval time
- Type of storage systems
- Physical distance between the storage systems
- Available bandwidth of the data link
- Input/output (I/O) load on the storage systems

To manage data exposure, you can change the consistency group interval time. The *consistency group time interval* specifies how often a Global Mirror and Metro Global Mirror session attempts to form a consistency group. When you reduce this value, it might be possible to reduce the data exposure of the session. A lower value causes the session to attempt to create consistency groups more frequently, which might also increase the processing load and message-traffic load on the storage systems.

## Data replication

Tivoli Storage Productivity Center for Replication supports several methods of data replication. The type of data replication that is associated with a given session is known as the *session type* (also known as a copy type).

The following table describes the session types that are supported by Tivoli Storage Productivity Center for Replication. Depending on the edition of Tivoli Storage Productivity Center for Replication that you are using, some of these session types might not be available.

Table 36. Session type summary

| Copy type       | Supported Software                                                                                                                                | Supported storage systems                                                                                                                                              | Description                                                                                                                                                                                                                                                                                |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Basic HyperSwap | Tivoli Storage Productivity Center for Replication Basic Edition for System z and Tivoli Storage Productivity Center for Replication for System z | <ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> </ul> | Basic HyperSwap replication is a special Metro Mirror replication method designed to provide high availability in the case of a disk storage system failure. Using Basic HyperSwap with Metro Mirror, you can configure and manage your synchronous Peer-to-Peer Remote Copy (PPRC) pairs. |

Table 36. Session type summary (continued)

| Copy type                     | Supported Software                                                              | Supported storage systems                                                                                                                                                                                                                                           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------------------|---------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FlashCopy                     | Tivoli Storage Productivity Center all editions                                 | <ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> </ul> | <p>FlashCopy replication creates a point-in-time copy in which the target volume contains a copy of the data that was on the source volume when the FlashCopy was established. Using FlashCopy, your data exists on the second set of volumes in the same storage system, and can be restored to the first set of volumes.</p> <p>SAN Volume Controller or Storwize V7000 FlashCopy sessions are managed by using FlashCopy consistency groups. Sessions for IBM TotalStorage Enterprise Storage Server (ESS) and IBM DS6000 and DS8000 are not managed by using FlashCopy consistency groups.</p> |
| Snapshot                      | Tivoli Storage Productivity Center all editions                                 | The XIV system                                                                                                                                                                                                                                                      | <p>Snapshot is a session type that creates a point-in-time copy of a volume or set of volumes without having to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.</p>                                                                                                                                                                                                                                                                                                                                              |
| Metro Mirror Single Direction | Tivoli Storage Productivity Center for Replication Two Site Business Continuity | <ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> </ul> | <p>Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. The source is located in one storage system and the target is located in another storage system. Using Metro Mirror, your data exists on the second site that is less than 300 KM away, and can be restored to the first site.</p>                                                                                                                                                                                                                                |

Table 36. Session type summary (continued)

| Copy type                                    | Supported Software                                                              | Supported storage systems                                                                                                                                                                                                                                                                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------------------------|---------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Metro Mirror Failover/Failback               | Tivoli Storage Productivity Center for Replication Two Site Business Continuity | <ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> <li>• The XIV system</li> </ul> | <p>Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. Using Metro Mirror Failover / Failback, your data exists on the second site that is less than 300 KM away. You can use failover and failback to switch the direction of the data flow. This ability enables you to run your business from the secondary site.</p> <p>Using Metro Mirror with HyperSwap, your data exists on the second site that is less than 300 KM away. The data can be restored to the first site. You can also use failover for a backup copy of the data if your primary volumes encounter a permanent I/O error.</p> |
| Metro Mirror Failover/Failback with Practice | Tivoli Storage Productivity Center for Replication Two Site Business Continuity | <ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> </ul>                           | <p>Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. The source is located in one storage system and the target is located in another storage system. Metro Mirror Failover / Failback with Practice combines Metro Mirror and FlashCopy to provide a point-in-time copy of the data on the remote site.</p>                                                                                                                                                                                                                                                                                     |

Table 36. Session type summary (continued)

| Copy type                                             | Supported Software                                                              | Supported storage systems                                                                                                                                                                                                                                           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------------------------------------|---------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Global Mirror Single Direction                        | Tivoli Storage Productivity Center for Replication Two Site Business Continuity | <ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> </ul> | Global Mirror is a method of asynchronous, remote data replication between two sites that are over 300 KM apart. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system. Using Global Mirror, your data exists on the second site that is more than 300 KM away, and can be restored to the first site.                                                                              |
| Global Mirror Either Direction with Two-Site Practice | Tivoli Storage Productivity Center for Replication Two Site Business Continuity | <ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> </ul>                                                                                              | Global Mirror is a method of asynchronous, remote data replication between two sites that are over 300 KM apart. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system. Global Mirror Either Direction with Two-Site Practice combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on either the primary or secondary sites that are over 300 KM apart. |

Table 36. Session type summary (continued)

| Copy type                                     | Supported Software                                                              | Supported storage systems                                                                                                                                                                                                                                                                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------------------|---------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Global Mirror Failover/Failback               | Tivoli Storage Productivity Center for Replication Two Site Business Continuity | <ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> <li>• The XIV system</li> </ul> | <p>Global Mirror is a method of asynchronous, remote data replication between two sites that are over 300 KM apart. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system. Using Global Mirror Failover / Failback, your data exists on the second site that is more than 300 KM away, and you can use failover and failback to switch the direction of the data flow. This ability enables you to run your business from the secondary site.</p> |
| Global Mirror Failover/Failback with Practice | Tivoli Storage Productivity Center for Replication Two Site Business Continuity | <ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS8000</li> <li>• System Storage DS6000</li> <li>• SAN Volume Controller</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> </ul>                           | <p>Global Mirror is a method of asynchronous, remote data replication between two sites that are over 300 KM apart. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system. Global Mirror Failover / Failback with Practice combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on a remote site at a distance over 300 KM away from your first site.</p>                                                             |

Table 36. Session type summary (continued)

| Copy type           | Supported Software                                                                | Supported storage systems                                                                                                                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------|-----------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Metro Global Mirror | Tivoli Storage Productivity Center for Replication Three Site Business Continuity | <ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800 (only H1 site)</li> <li>• System Storage DS8000</li> </ul> | <p>Metro Global Mirror is a method of continuous, remote data replication that operates between three sites that varying distances apart. Metro Global Mirror combines Metro Mirror synchronous copy and Global Mirror asynchronous copy into a single session, where the Metro Mirror target is the Global Mirror source. Using Metro Global Mirror and Metro Global Mirror with HyperSwap, your data exists on a second site that is less than 300 KM away, and a third site that is more than 300 KM away. Metro Global Mirror uses both Metro Mirror and Global Mirror Failover / Failback to switch the direction of the data flow. This ability enables you to run your business from the secondary or tertiary sites.</p> <p>Using Basic HyperSwap with Metro Global Mirror, you can configure and manage three-site continuous replication needed in a disaster recovery event.</p> |

Table 36. Session type summary (continued)

| Copy type                         | Supported Software                                                                | Supported storage systems                                                                                                                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------|-----------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Metro Global Mirror with Practice | Tivoli Storage Productivity Center for Replication Three Site Business Continuity | <ul style="list-style-type: none"> <li>TotalStorage Enterprise Storage Server Model 800 (only H1 site)</li> <li>System Storage DS8000</li> </ul> | Using Metro Global Mirror with Practice, you can practice your disaster recovery actions while maintaining disaster recovery capabilities. Your data exists on a second site that is less than 300 KM away, and a third site that is more than 300 KM away. Metro Global Mirror uses both Metro Mirror and Global Mirror Failover / Failback to switch the direction of the data flow; as a result, you can run your business from the secondary or tertiary sites, and simulate a disaster. |

### Basic HyperSwap (ESS, DS6000, and DS8000)

Basic HyperSwap is an entitled copy services solution for z/OS version 1.9 and later. It provides high availability of data in the case of a disk storage system failure. Basic HyperSwap is not a disaster recovery solution. If a session is suspended but the suspend was not caused by a HyperSwap trigger, no freeze is done to ensure consistency of the session.

When HyperSwap is combined with Metro Mirror and Metro Global Mirror replication, you can prepare for disaster recovery and ensure high availability. If a session is suspended but the suspend was not caused by a HyperSwap trigger, a freeze is done to ensure consistency of the session.

**Note:** This replication method is available on only ESS, DS6000, and DS8000 storage systems, and on management servers running IBM Tivoli Storage Productivity Center for Replication Basic Edition for System z or IBM Tivoli Storage Productivity Center for Replication for System z.

Basic HyperSwap replication performs the following actions:

- Manages CKD volumes in Metro Mirror (synchronous peer-to-peer remote copy [PPRC]) relationships.
- Permits only CKD volumes to be added to the HyperSwap session. The graphical user interface (GUI) shows only CKD volumes when you add a copy set. The command-line interface (CLI) fails to add a copy set if a fixed block volume is specified.
- Monitoring for events that indicate a storage device has failed.
- Determining whether the failing storage device is part of a Metro Mirror (synchronous PPRC) pair.
- Determining the action to be taken from policy.
- Ensuring that data consistency is not violated.



- Swapping the I/O between the primary logical devices in the consistency group with the secondary logical devices in the consistency group. A swap can be from the preferred logical devices to the alternate logical devices or from the alternate logical devices to the preferred logical devices.

### **Metro Mirror Failover/Failback with HyperSwap**

Metro Mirror Failover/Failback uses HyperSwap to configure and manage synchronous Peer-to-Peer Remote Copy (PPRC) pairs.

### **Metro Global Mirror with HyperSwap**

Metro Global Mirror with HyperSwap is a z/OS replication feature that provides the three-site continuous replication needed in a disaster recovery event.

**Important:** If HyperSwap occurs by event when running a Metro Global Mirror with a HyperSwap session, a full copy of the data occurs to return to a full three-site configuration. If you issue a HyperSwap command when running a Metro Global Mirror with a HyperSwap session, a full copy does not occur. A full copy is required only for an unplanned HyperSwap or a HyperSwap initiated using the z/OS SETHS SWAP command.

### **Example**

Jane is using multiple DS8000 storage systems. Her host applications run on z/OS and her z/OS environment has connectivity to the DS8000 storage systems. She has a site in Manhattan and a bunker in Hoboken. While she does not need a disaster recovery solution, she does need a high-availability solution to keep her applications running around the clock. Jane is worried that if a volume fails on her DS8000 in Manhattan, her database application will not be able to operate. Even a small downtime can cost Jane thousands of dollars. Jane uses a Basic HyperSwap session to mirror her data in Manhattan to her secondary DS8000 in Hoboken. If a volume at the Manhattan site fails, Basic HyperSwap automatically directs application I/O to the mirrored volumes in Hoboken.

### **FlashCopy**


FlashCopy replication creates a point-in-time copy in which the target volume contains a copy of the data that was on the source volume when the FlashCopy was established.

The ESS, DS6000, and DS8000 platforms provide multiple logical subsystems (LSSs) within a single physical subsystem, while the following platforms provide multiple I/O groups:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

All platforms support local replication in which the source volume is located in one LSS or I/O group and the target volume is located in the same or another LSS or I/O group. Using FlashCopy, you can reference and update the source volume and target volume independently.

The following figure illustrates how a FlashCopy session works.



### Example

Jane uses FlashCopy to make a point-in-time copy of the customer data in existing international accounts. Every night, the bank's servers perform batch processing. Jane uses FlashCopy to create checkpoint restarts for the batch processing in case the batch processing fails. In her batch processing, the first step is to balance all international accounts, with a FlashCopy taken of the resulting data. The second step in her batch processing is to process the international disbursements. If the second step in the batch process fails, Jane can use the FlashCopy made of the first step to repeat the second step, instead of beginning the entire process again. Jane also writes a CLI script that performs a FlashCopy every night at 11:59 PM, and another script that quiesces the database. She backs this data up on tape on her target storage system, and then sends the tape to the bank's data facility in Oregon for storage.

### Snapshot

Snapshot is a session type that creates a point-in-time copy of a volume or set of volumes. You do not have to define a specific target volume. The target volumes of a Snapshot session are automatically created when the snapshot is created.

The XIV system uses advanced snapshot architecture to create a large number of volume copies without affecting performance. By using the snapshot function to create a point-in-time copy, and to manage the copy, you can save storage. With the XIV system snapshots, no storage capacity is used by the snapshot until the source volume (or the snapshot) is changed.

The following figure illustrates how a Snapshot session works.



### Example

Jane's host applications are using an XIV system for their back-end storage. With the XIV system, Jane can create a large number of point-in-time copies of her data. The snapshot function ensures that if data becomes corrupted, she can restore the data to any number of points in time.

Jane sets up a Snapshot session by using Tivoli Storage Productivity Center for Replication and specifies the volumes on the XIV system that are used by her host applications. Jane does not have to provision target volumes for all the snapshots she intends to make. She can quickly get a single Snapshot session configured and ready.

When the session is configured, Jane writes a CLI script that performs a Create Snapshot command to the session every two hours. If a problem occurs, such as data becoming corrupted, Jane can find a snapshot of the data before the problem occurred. She can restore the data to that point.

By creating a set of snapshots of the data, Jane can also schedule batch processing against that data every day. She can use the batch processing to analyze certain trends in the market without causing any effect to the host applications.

### Metro Mirror

Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart. The source is located in one storage system and the target is located in another storage system.

**Attention:** For Tivoli Storage Productivity Center for Replication for System z sessions containing Metro Mirror relationships, ensure that the session does not contain system volumes (such as paging volumes) unless the session is enabled for HyperSwap. If HyperSwap is not enabled, a freeze that is issued through Tivoli Storage Productivity Center for Replication might cause Tivoli Storage Productivity Center for Replication processing to freeze. This situation might prevent the session from ensuring that the data is consistent.


Metro Mirror replication maintains identical data in both the source and target. When a write is issued to the source copy, the changes made to the source data are propagated to the target before the write finishes posting. If the storage system goes down, Metro Mirror provides zero loss if data must be used from the recovery site.

A Metro Mirror session in Global Copy mode creates an asynchronous relationship to accommodate the high volume of data migration. As a result, the data on the target system might no longer be consistent with the source system. The Metro Mirror session switches back to a synchronous relationship when a Metro Mirror Start command is reissued. In addition, you can start a Metro Mirror session in Global Copy mode and toggle between Metro Mirror and Global Copy modes to accommodate time periods in which you value host input/output (I/O) response time over data consistency.

**Tip:** To determine if there is any out of synch data that must to be copied over before the session is consistent, check the percent that is complete in the session details panel.

### Metro Mirror Single Direction


The following figure illustrates how a Metro Mirror Single Direction session works.



## Metro Mirror Failover/Failback

Using Metro Mirror with failover/failback, your data exists on the second site that is less than 300 KM away, and you can use failover/failback to switch the direction of the data flow. This session type enables you to run your business from the secondary site, and to copy changes made at the second site back to the primary site when you want to resume production at the primary site.

The following figure illustrates how a Metro Mirror with Failover/Failback session works.




## Metro Mirror Failover/Failback with Practice

A Metro Mirror Failover/Failback with Practice session combines Metro Mirror and FlashCopy to provide a point-in-time copy of the data on the remote site. You can use this session type to practice what you might do if a disaster occurred, without losing your disaster recovery capability.

This solution consists of two host volumes and an intermediate volume.

The following figure illustrates how a Metro Mirror Failover/Failback with Practice session works.



## Metro Mirror Failover/Failback with HyperSwap

A Metro Mirror Failover/Failback session can be enabled to have HyperSwap capabilities. To enable HyperSwap the following circumstances must apply:

- The session is running on a Tivoli Storage Productivity Center for Replication server running on IBM z/OS.
- The volumes are only for TotalStorage Enterprise Storage Server, System Storage DS8000, and DS6000 systems.
- The volumes are count key data (CKD) volumes that are attached to the z/OS system.


Metro Mirror Failover/Failback with HyperSwap combines the high availability of Basic HyperSwap with the redundancy of a two-site Metro Mirror

Failover/Failback solution when managing count key data (CKD) volumes on z/OS. If the primary volumes encounter a permanent I/O error, the I/O is automatically swapped to the secondary site with little to no impact on the application.

A swap can be planned or unplanned. A planned swap occurs when you issue a HyperSwap command from the **Select Action** list in the graphical user interface (GUI) or when you issue a **cmdsess -action hyperswap** command.

The following figure illustrates how a Metro Mirror Failover/Failback session enabled for HyperSwap works.

#### Metro Mirror Failover/Failback



For more information about enabling HyperSwap, see “Managing a session with HyperSwap and Open HyperSwap replication” on page 100.

#### Metro Mirror Failover/Failback with Open HyperSwap

A Metro Mirror Failover/Failback session can be enabled to have Open HyperSwap capabilities. To enable Open HyperSwap the following circumstances must apply:


- The volumes in the session are System Storage DS8000 5.1 or later volumes.
- The volumes in the session are fixed block and are mounted to IBM AIX 5.3 or AIX 6.1 hosts with the following modules installed:
  - Subsystem Device Driver Path Control Module (SDDPCM) version 3.0.0.0 or later
  - Multi-Path Input/Output (MPIO) module (the version that is provided with AIX version 5.3 or 6.1)
- The connections between the AIX host systems and the Tivoli Storage Productivity Center for Replication server have been established.

Metro Mirror Failover/Failback with Open HyperSwap combines the high availability of Basic HyperSwap on z/OS for fixed block AIX volumes with the redundancy of a two-site Metro Mirror Failover/Failback solution. If the primary volumes encounter a permanent I/O error, the I/O is automatically swapped to the secondary site with little to no impact on the application.

A swap can be planned or unplanned. A planned swap occurs when you issue a HyperSwap command from the **Select Action** list in the GUI or when you issue a **cmdsess -action hyperswap** command.

The following figure illustrates how a Metro Mirror Failover/Failback session enabled for Open HyperSwap works.

## Metro Mirror Failover/Failback



For more information about enabling Open HyperSwap, see “Managing a session with HyperSwap and Open HyperSwap replication” on page 100.

## Examples

### Metro Mirror Single Direction

At the beginning of a work week, Jane is notified that between 10:00 AM and 11:00 AM on the next Friday, power in her building is going to be shut off. Jane does not want to lose any transactions during the power outage, so she decides to transfer operations to the backup site during the outage. She wants a synchronous copy method with no data loss for her critical business functions, so she chooses Metro Mirror, which can be used between locations that are less than 300 KM apart.

In a synchronous copy method, when a write is issued to change the source, the change is propagated to the target before the write is completely posted. This method of replication maintains identical data in both the source and target. The advantage of this is when a disaster occurs, there is no data loss at the recovery site because both writes must complete before signaling completion of a write to the source application. Because the data must be copied to both System Storage DS8000 devices before the write is completed, Jane can be sure that her data is safe.

The night before the planned outage, Jane quiesces her database and servers in San Francisco and starts the database and servers in Oakland. To accomplish this task, Jane issues the Suspend and Recover commands, and then issues the Start command on the secondary site. She powers off her equipment in San Francisco to avoid any power spikes during reboot after the power is turned back on.

### Metro Mirror in Global Copy mode

At the beginning of a work week, Jane is notified that between 10:00 AM and 11:00 AM on the next Friday, power in her building is going to be shut off. Jane does not want to lose any transactions during the power outage, so she decides to transfer operations to their backup site during the outage. She wants a synchronous copy method with no data loss for her critical business functions, so she chooses Metro Mirror, which can be used between locations that are less than 300 KM apart.

Jane wants to limit her application impact while completing the initial Metro Mirror synchronization, so she begins her session in Global Copy mode. After she sees that about 70% of the data has been copied, Jane decides to switch the session into Metro Mirror mode, assuring data consistency.

### Metro Mirror with Practice

Jane wants to run a Metro Mirror with Practice from San Francisco to Oakland. She wants to verify her recovery procedure for the Oakland site,

but she cannot afford to stop running her Metro Mirror session while she takes time to practice a recovery. By using a Metro Mirror with Practice session, Jane is able to practice her disaster recovery scenario in Oakland while her Metro Mirror session runs uninterrupted. By practicing running her applications at the Oakland site, Jane is better prepared to make a recovery if a disaster ever strikes the San Francisco site.

While her session is running in a Prepared state, Jane practices a recovery at her Oakland site by issuing the Flash command. This command momentarily pauses the session and starts a FlashCopy to the H2 volumes. As soon as the FlashCopy is started, her session will be restarted. These FlashCopy files create a consistent version of the data on the H2 volume that she can use for recovery testing, while her session continues to replicate data from San Francisco to Oakland. As a result, she can carry out her recovery testing without stopping her replication for any extended duration of time.

If at some point in time, the Metro Mirror session suspends due to a failure, Jane can use the practice session to restart her data replication while maintaining a consistent copy of the data at the Oakland site, in case of a failure during the resynchronization process. When the session is suspended, she can issue a Recover command to create a consistent version of the data on the H2 volumes. After the Recover command completes, she can issue the Start H1->H2 command to resynchronize the data from the San Francisco site to the Oakland site. If a failure occurs before her restarted session is in the Prepared state, she has a consistent version of the data on the H2 volumes. She must simply issue the Recover command to put the session into Target Available state and make the H2 volumes accessible from her servers. If the session was not in the Prepared state when it suspended, the subsequent Recover command does not issue the FlashCopy files to put the data on the H2 volumes. This means that the consistent data on the H2 volumes are not overwritten if the data to be copied to them is not consistent.

### **Metro Mirror Failover/Failback enabled for Open HyperSwap**

Jane wants to run a Metro Mirror with Practice from San Francisco to Oakland. She wants to verify her recovery procedure for the Oakland site, but she cannot afford to stop running her Metro Mirror session while she takes time to practice a recovery. By using a Metro Mirror with Practice session, Jane is able to practice her disaster recovery scenario in Oakland while her Metro Mirror session runs uninterrupted. By practicing running her applications at the Oakland site, Jane is better prepared to make a recovery if a disaster ever strikes the San Francisco site.

While her session is running in a Prepared state, Jane practices a recovery at her Oakland site by issuing the Flash command. This command momentarily pauses the session and starts a FlashCopy to the H2 volumes. As soon as the FlashCopy is started, the session is restarted. These FlashCopy files create a consistent version of the data on the H2 volume that she can use for recovery testing, while her session continues to replicate data from San Francisco to Oakland. As a result, she can carry out her recovery testing without stopping her replication for any extended duration of time.

If the Metro Mirror session suspends due to a failure, Jane can use the practice session to restart her data replication while maintaining a consistent copy of the data at the Oakland site, in case of a failure during the resynchronization process. When the session is suspended, she can

issue a Recover command to create a consistent version of the data on the H2 volumes. After the Recover command completes, she can issue the Start H1->H2 command to resynchronize the data from the San Francisco site to the Oakland site. If a failure occurs before her restarted session is in the Prepared state, she has a consistent version of the data on the H2 volumes. She must simply issue the Recover command to put her session into Target Available state and make the H2 volumes accessible from her servers. If the session was not in the Prepared state when it suspended, the subsequent Recover command does not issue the FlashCopy files to put the data on the H2 volumes. This means that the consistent data on the H2 volumes are not overwritten if the data to be copied to the volumes is not consistent.

### Selecting a HyperSwap session

A global insurance company has elected to use Tivoli Storage Productivity Center for Replication to manage its disaster recovery environment. Jane wants minimal data exposure, both for planned outages such as routine maintenance, and for unplanned disasters. They have CKD volumes on System Storage DS8000 devices, and use z/OS mainframes. They have two data centers in New York.

Jane reviews the Tivoli Storage Productivity Center for Replication documentation, and chooses a Metro Mirror recovery solution, since her priority is being protected against regional disasters. Jane decides to use Metro Mirror solution, because her company has two data centers located near each other. Jane realizes that because she uses z/OS, CKD, and System Storage DS8000 hardware, she is also able to use a HyperSwap solution. Using Metro Mirror Failover/Failback with HyperSwap, Jane can minimize application impact, while maintaining seamless failover to her secondary site. Jane decides Metro Mirror Failover/Failback with HyperSwap is best for the needs of her company.

After installing and configuring Tivoli Storage Productivity Center for Replication on z/OS, Jane starts the Tivoli Storage Productivity Center for Replication GUI. She adds the Tivoli Storage Productivity Center for Replication storage devices she intends to use on all three sites. From the Session Overview panel, Jane launches the Create Session wizard, and selects the **Metro Mirror Failover/Failback** session type. Then, she continues the wizard, she selects the **Manage H1H2 with HyperSwap** option. After finishing the wizard, Jane clicks **Launch Add Copy Sets Wizard**. She completes this wizard, and issues a **Start H1->H2** command. After the initial copy is completed, Jane is safely replicating her data between both sites. She can also issue a HyperSwap between sites 1 and 2, enabling her to switch sites with minimal application impact during either a disaster or maintenance period.

### Performing a planned HyperSwap

Jane's company has successfully been using Metro Mirror Failover/Failback with HyperSwap sessions for the past three months. However, Jane needs to perform maintenance on an H1 box. During this time, Jane does not want her applications or replication to be interrupted. To prevent this from happening, shortly before the maintenance is scheduled to begin, Jane decides to use the Tivoli Storage Productivity Center for Replication GUI to perform a HyperSwap to the H2 volumes. This transitions the applications so that they write to H2. To perform a planned HyperSwap, Jane issues a HyperSwap command.



### Understanding what happens when an unplanned HyperSwap occurs

Several weeks after the planned maintenance at Jane's company is completed, an incident occurs at the H1 site. A disk controller fails, causing one of the H1 volumes to encounter a permanent I/O error. Fortunately, Jane's data is safe because she used Metro Mirror Failover/Failback with HyperSwap, and her H2 volume is an exact duplicate of the H1 volume. When the permanent I/O error is detected, a HyperSwap is triggered. The application seamlessly transitions to writing to the H2 volumes. Her data is safe, and her applications are not interrupted.

Jane configured an Simple Network Management Protocol (SNMP) listener to alert her to any events, so she receives the SNMP event indicating that a HyperSwap occurred. Jane investigates the cause of the HyperSwap and uses the z/OS console to identify the volume that triggered the HyperSwap. Jane replaces the faulty disk controller. Then, to recover from the unplanned HyperSwap, Jane issues the **Start H2->H1** command.

### Global Mirror


Global mirror is a method of asynchronous, remote data replication between two sites that are over 300 kilometers (km) apart. It maintains identical data in both the source and target, where the source is located in one storage system and the target is located in another storage system.

The data on the target is typically written a few seconds after the data is written to the source volumes. When a write is issued to the source copy, the change is propagated to the target copy, but subsequent changes are allowed to the source before the target verifies that it has received the change. Because consistent copies of data are formed on the secondary site at set intervals, data loss is determined by the amount of time since the last consistency group was formed. If your system stops, Global Mirror might lose some data that was being transmitted when the disaster occurred. Global Mirror still provides data consistency and data recoverability in the event of a disaster.

### Global Mirror Single Direction

A Global Mirror single direction session allows you to run your Global Mirror replication from only the primary site.

For ESS, DS6000, and DS8000 storage systems, each copy set in the Global Mirror single direction session consists of two host volumes and a journal volume. The following figure illustrates how a Global Mirror single direction session works on an ESS, DS6000, and DS8000 storage systems:



For the following storage systems, each copy set in the Global Mirror Single Direction session consists of two host volumes:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

The following figure illustrates how a Global Mirror Single Direction session works on these storage systems:




### Global Mirror Either Direction with Two-Site Practice (ESS, DS6000, and DS8000)

A Global Mirror Either Direction with Two Site Practice session allows you to run Global Mirror replication from either the primary or secondary site. It combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on a remote site at a distance over 300 km away from your first site. This practice session allows you to create practice volumes on both the primary and secondary site to practice what you might do if a disaster occurred, without losing your disaster recovery capability.

**Note:** This replication method is available on only ESS, DS6000, and DS8000 storage systems.


The session consists of two host volumes, two intermediate volumes, and two journal volumes. The following figure illustrates how a Global Mirror either direction with two-site practice session works:



## Global Mirror Failover/Failback

Using Global Mirror failover/failback, your data exists on the second site that is more than 300 km away, and you can use failover/failback to switch the direction of the data flow. This enables you to run your business from the secondary site.

For ESS, DS6000, and DS8000 storage systems, each copy set in the Global Mirror failover/failback session consists of two host volumes and a journal volume. The following figure illustrates how a Global Mirror failover/failback session works on an ESS, DS6000, or DS8000 storage systems:



For the following storage systems, each copy set in the Global Mirror failover/failback session consists of two host volumes:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified
- The XIV system


The following figure illustrates how a Global Mirror failover/failback session works on these storage systems.



## Global Mirror Failover/Failback with Practice

A Global Mirror failover/failback with practice session combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on a remote site at a distance over 300 km away from your primary site. You can use this to practice what you might do if a disaster occurred, without losing your disaster recovery capability.


For ESS, DS6000, and DS8000 storage systems, each copy set in the Global Mirror failover/failback with practice session consists of two host volumes, an intermediate volume, and a journal volume. The following figure illustrates how a Global Mirror failover/failback with practice session works on an ESS, DS6000, or DS8000 storage system:



For the following storage systems, each copy set in the Global Mirror failover/failback with Practice session consists of two host volumes and an intermediate volume:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

The following figure illustrates how a Global Mirror Failover/Failback with Practice session works on these storage systems:



## Examples

### Global Mirror Single Direction

Although Jane's FlashCopy and Metro Mirror copies were both planned, Jane realizes that sometimes unforeseen things happen, and she wants to make sure her data is safe. Because Jane works in San Francisco, she wants her other site to be far away in case of a localized disaster. Her other site is based in Houston. Jane's foresight pays off when a minor earthquake occurs in San Francisco and power and communications both go down. Fortunately, Jane has arranged for the data on customer accounts that have recently opened or closed to be asynchronously copied in Houston, using

Global Mirror. Jane risks losing the bytes of data that were being processed when the tremor disrupted the San Francisco process, but she views that as a minor inconvenience when weighed next to the value of backing up her data in a non-earthquake zone.

#### **Global Mirror with Practice**

Jane wants to run a Global Mirror with practice from San Francisco to Houston. She wants to verify her recovery procedure for the Houston site, but she cannot afford to stop running her Global Mirror session while she takes time to practice a recovery. By using a Global Mirror with Practice session, Jane is able to practice her disaster recovery scenario in Houston while her Global Mirror session runs uninterrupted. By practicing running her applications at the Houston site, Jane will be better prepared to make a recovery if a disaster ever strikes the San Francisco site.

#### **Global Mirror Either Direction with Two-Site Practice**

Jane wants to run a Global Mirror with practice from San Francisco to Houston. She wants to verify her recovery procedure for the Houston site, but she cannot afford to stop running her Global Mirror session while she takes time to practice a recovery. By using a Global Mirror either direction two-site with practice session, Jane is able to practice her disaster recovery scenario in Houston while her Global Mirror session runs uninterrupted. By practicing running her applications at the Houston site, Jane will be better prepared to make a recovery if a disaster ever strikes the San Francisco site.

Jane can use the Global Mirror either direction with two-site practice session to run asynchronous consistent data replication from either the San Francisco site or the Houston site. (She can practice her disaster recovery at the target site, no matter what her current production site is.) Jane's business is able to run a consistent Global Mirror session from its Houston site back to San Francisco while running a production at the Houston site.

### **Setting up Global Mirror for Resource Groups on System Storage DS8000**

If resource groups are defined on a System Storage DS8000, Global Mirror session IDs might be defined for some users. Tivoli Storage Productivity Center for Replication does not automatically determine which session IDs are valid. To determine which session IDs are valid, you must modify the `rmserver.properties` file and add the following property:

```
gm.master.sessionid.gm_role,session_name = xx
```

where *gm\_role* is the role that has the master volume (for example, H1 in a Global Mirror failover/failback session), *session\_name* is the name of the session that uses the session ID, and *xx* is the decimal number for the session ID.

**Important:** System Storage DS8000 represents session IDs as a two-digit hexadecimal number. Use the decimal version of that number. For example, if you want a Global Mirror failover/failback session to use a session ID of 0F, the decimal number is 15 as shown in the following example:

```
gm.master.sessionid.H2.11194_wprac=15
```

### **Metro Global Mirror (ESS 800 and DS8000)**

Metro Global Mirror is a method of continuous, remote data replication that operates between three sites that varying distances apart. Metro Global Mirror

combines Metro Mirror synchronous copy and Global Mirror asynchronous copy into a single session, where the Metro Mirror target is the Global Mirror source.

**Notes:**

- This replication method is available on only ESS 800 and DS8000 storage systems.
- You can select ESS 800 storage systems in only the H1 volume role. All other volume roles must use DS8000 volumes.
- You can mix ESS 800 and DS8000 volumes in the H1 volume role. If ESS 800 and DS8000 storage systems are both used in the H1 role, the DS8000 storage system performs Incremental Resync (IR), and the ESS 800 storage system performs a full copy. Because ESS 800 does not support the IR function, a full copy is required when changing from H1->H2->H3 to H1>H3 and from H2->H1->H3 to H2->H3.

Metro Global Mirror maintains a consistent copy of data at the remote site, with minimal impact to applications at the local site. This remote mirroring function works in combination with FlashCopy to meet the requirements of a disaster-recovery solution by providing the following features:

- Fast failover and failback
- Rapid reestablishment of three-site mirroring, without production outages
- Data currency at the remote site with minimal lag behind at the local site, an average of only 3 - 5 seconds for many environments
- Quick resynchronization of mirrored sites using only incremental changes

If IBM Tivoli Storage Productivity Center for Replication is running on z/OS, you can configure a Metro Global Mirror session to control the Metro Mirror relationship between the primary and secondary site using HyperSwap. With HyperSwap enabled, a failure on the primary storage system causes an automatic HyperSwap, transparently redirecting application I/O to the auxiliary storage system. The Global Mirror relationship continues to run uninterrupted throughout this process. With this configuration, you can achieve near zero data loss at larger distances.

Using synchronous mirroring, you can switch from local site H1 to remote site H2 during a planned or unplanned outage. It also provides continuous disaster recovery protection of site H2 through site H3, without the necessity of additional reconfiguration, if a switch from site H1 occurs. With this configuration, you can reestablish H2->H1->H3 recoverability while production continues to run at site H2. Additionally, this cascaded setup can reduce the load on-site H1 as compared to some multi-target (non-cascaded) three-site mirroring environments.


**Important:**

- If HyperSwap occurs by event when running a Metro Global Mirror with a HyperSwap session, a full copy of the data occurs to return to a full three-site configuration. If you issue a HyperSwap command when running a Metro Global Mirror with a HyperSwap session, a full copy does not occur. A full copy is required only for an unplanned HyperSwap or a HyperSwap initiated using the z/OS SETHS SWAP command.
- In Metro Global Mirror and Metro Global Mirror with Practice sessions, when the H1 is on an ESS 800, you might risk filling up the space efficient journal volumes. Because incremental resynchronization is not supported on the ESS 800, full copies are performed in many of the transitions.

## Metro Global Mirror

A Metro Global Mirror session with Practice combines Metro Mirror, Global Mirror, and FlashCopy across three sites to provide a point-in-time copy of the data on the third site.

The following figure illustrates how a Metro Global Mirror session works.




## Metro Global Mirror with Practice

A Metro Global Mirror session with Practice combines Metro Mirror, Global Mirror, and FlashCopy across three sites to provide a point-in-time copy of the data on the third site. You can use this session to practice what you might do if a disaster occurred without losing your disaster recovery capability.

The session consists of three host volumes, an intermediate volume, and a journal volume. The following figure illustrates how a Metro Global Mirror with Practice session works.

**Note:** A Metro Global Mirror with practice session can be created when the three-site license has been applied to the server.



**Note:** In Metro Global Mirror and Metro Global Mirror with Practice sessions, when the H1 is on an ESS 800, you might risk filling up the space efficient journal volumes. Because incremental resynchronization is not supported on the ESS 800, full copies are performed in many of the transitions.

## Examples

### Metro Global Mirror

Although Jane works in San Francisco, she wants to give herself the ability to run her business from either Oakland (her secondary site) or Houston (her tertiary site). Jane can use Metro Global Mirror with failover/failback to switch the direction of the data flow, so that she can run her business from either Oakland or Houston. Metro Global Mirror means that Jane has zero data loss backup at her secondary site, and minimal data loss at her tertiary site.

### Metro Global Mirror with Practice

Jane wants to run a Metro Global Mirror with Practice from San Francisco to Houston. She wants to verify her recovery procedure for the Houston site. However, she cannot afford to stop running her Metro Global Mirror session while she takes time to practice a recovery. By using a Metro Global Mirror with Practice session, Jane is able to practice her disaster recovery scenario in Houston while her Metro Global Mirror session runs uninterrupted. By practicing running her applications at the Houston site, and being prepared to run her applications at the Oakland site if necessary, Jane will be better prepared to make a recovery if a disaster ever strikes the San Francisco site.

Jane can use Metro Global Mirror with Practice to switch the direction of the data flow, so that she can run her business from either Oakland or Houston. Using Metro Global Mirror, Jane has zero data loss backup at her secondary site; and minimal data loss at her tertiary site.

## Managing a session with HyperSwap and Open HyperSwap replication

HyperSwap and Open HyperSwap provide high availability of data in the case of a primary disk storage system failure. When a failure occurs in writing input/output (I/O) to the primary storage system, the failure is detected by IOS, and IOS automatically swaps the I/O to the secondary site with no user interaction and little or no application impact.

### Sessions that can be enabled for HyperSwap or Open HyperSwap:

You can create sessions that have HyperSwap or Open HyperSwap capabilities. Enabling swapping provides a session with a highly available business continuity solution.

### Sessions that can enable HyperSwap

The following session types can enable HyperSwap:

- Basic HyperSwap
- Metro Mirror with Failover/Failback
- Metro Global Mirror

To enable HyperSwap, the following circumstances must apply:

- The session is running on an Tivoli Storage Productivity Center for Replication server that is running on IBM z/OS.



- The volumes are only for TotalStorage Enterprise Storage Server, System Storage DS8000, and DS6000 systems.
- The volumes are count key data (CKD) volumes that are attached to the z/OS system.

### **Sessions that can enable Open HyperSwap**

Only Metro Mirror with Failover/Failback session type can enable Open HyperSwap.

To enable Open HyperSwap, the following circumstances must apply:

- The volumes in the session are only for System Storage DS8000 5.1 or later volumes.
- The volumes in the session are fixed block and mounted to IBM AIX 5.3 or AIX 6.1 hosts with the following modules installed:
  - Subsystem Device Driver Path Control Module (SDDPCM) version 3.0.0.0 or later
  - Multi-Path Input/Output (MPIO) module (the version that is provided with AIX version 5.3 or 6.1)
- The connections between the AIX host systems and the Tivoli Storage Productivity Center for Replication server have been established.

### **Setting up the environment for HyperSwap:**

You must set up an environment that supports HyperSwap before attempting to enable HyperSwap for a IBM Tivoli Storage Productivity Center for Replication session.

The following steps must be completed before HyperSwap can be enabled. For more information about these steps, see the IBM Tivoli Storage Productivity Center for Replication for System z Installation and Configuration Guide

1. Install IBM Tivoli Storage Productivity Center for Replication for System z.
2. Perform the post installation tasks of setting up the data store and other necessary system settings.
3. Ensure that all RESERVEs are converted to global enqueues (ENQs).
4. Ensure that all volumes in the session that you are enabling for HyperSwap are attached to the IBM z/OS system that is running Tivoli Storage Productivity Center for Replication.

### **Setting up the environment for Open HyperSwap:**

You must set up an environment that supports Open HyperSwap before attempting to enable Open HyperSwap for a IBM Tivoli Storage Productivity Center for Replication session.

The following steps must be completed before Open HyperSwap can be enabled:

1. Ensure that the IBM AIX hosts and IBM System Storage DS8000 to meet the following hardware and software requirements:

#### **AIX requirements**

Open HyperSwap support requires AIX version 5.3 or 6.1. (You can find the supported AIX version for each Tivoli Storage Productivity Center for Replication release in the support matrix at <http://www-01.ibm.com/support/docview.wss?rs=40&context=SSBSEX&context=SSMN28&context=SSMMUP&context=SS8JB5>)

&context=SS8JFM&uid=swg21386446&loc=en\_US&cs=utf-8&lang=en.  
Click the link for the applicable release under Agents, Servers and GUI.)

You must have the following AIX modules installed:

- Subsystem Device Driver Path Control Module (SDDPCM) version 3.0.0.0 or later
- Multi-Path Input/Output (MPIO) module (the version that is provided with AIX version 5.3 or 6.1)

#### **System Storage DS8000 hardware requirements**

Only System Storage DS8000 storage systems are supported. Open HyperSwap requires System Storage DS8000 5.1 or later.

Open HyperSwap does not support High Availability Cluster Multi-Processing (HACMP).

2. Create connections from Tivoli Storage Productivity Center for Replication to the AIX hosts (see “Adding a host system connection” on page 666).
3. Assign copy set volumes from the storage device to the host using the System Storage DS8000 command-line interface (CLI) or graphical user interface (GUI).
4. Run the AIX `cfgmgr` command to discover the volumes assigned to the host.

#### **Considerations for Open HyperSwap and the AIX host:**

- A single session that has Open HyperSwap enabled can manage multiple hosts; however, each host can be associated with only one session. Multiple hosts can share the same session.
- For AIX 5.3, a single host can manage a maximum of 1024 devices that have been enabled for Open HyperSwap on the host, with 8 logical paths configured for each copy set in the session. For AIX 6.1, a single host can manage a maximum of 1024 devices that have been enabled for Open HyperSwap on the host, with 16 logical paths configured for each copy set in the session.
- If an application on the host has opened a device, a Tivoli Storage Productivity Center for Replication session for that device cannot be terminated. The **Terminate** command fails. To terminate the session, you must either close the application or remove the copy sets from the session. If you remove copy sets from the session, you must ensure that the application writes to the correct volume when the copy set relationship is restored.
- It is possible for Open HyperSwap to fail on a subset of hosts for the session and work on the remaining hosts for the same session. In this situation, you must determine the best action to take if the application is writing to volumes on the source system as well as volumes on the target system. Contact the IBM Support Center if you need assistance determining the best solution for this issue.
- To enable support for Open HyperSwap on the host, refer to the *IBM System Storage Multipath Subsystem Device Driver User's Guide*.

#### *Configuring timers to support Open HyperSwap:*

There are configurable timeout values for the storage system, IBM Tivoli Storage Productivity Center for Replication, and IBM AIX hosts systems that can affect the operation of Open HyperSwap.

The following list describes the various timeout values that can affect Open HyperSwap:

### Storage system quiesce timeout value

This quiesce timeout timer begins when the storage system starts a quiesce operation. When the timer value expires, input/output (I/O) is resumed on the primary device. The default timeout value is two minutes, but the value can be set from 30 to 600 seconds. To set the quiesce timeout value, see the information about the **chdev** command in the *IBM System Storage Multipath Subsystem Device Driver User's Guide*.

### Storage system long busy timeout value

This timeout value is the time in seconds that the logical subsystem (LSS) consistency group volume stays in the long busy state after a remote mirror and copy error is reported.

### Timeout values for the applications that are on the host

The various applications that are running on the host have timeout values. The timeout values vary depending on the application.

### Considerations for setting timers

Consider the following information for setting timers:

- If the host quiesce timer is set to a shorter value than the Tivoli Storage Productivity Center for Replication response timer, it is possible that an I/O swap failure can occur. If a storage system triggers an unplanned failover and if the storage system quiesce timer expires before Tivoli Storage Productivity Center for Replication responds, the host attempts to write I/O to the primary volume where the loss of access occurred. If the hardware condition that caused the loss of access continues, the attempt to write I/O fails again and an unplanned Open HyperSwap is not performed.
- If the host quiesce timer is set to a longer value than the Tivoli Storage Productivity Center for Replication response timer, an application timeout might occur if Open HyperSwap takes too long to complete.

### Enabling a Session for HyperSwap or Open HyperSwap:

Enabling HyperSwap or Open HyperSwap for a session provides a combined business recovery and business continuity solution.

To ensure that your environment supports HyperSwap or Open HyperSwap, see “Setting up the environment for HyperSwap” on page 101 or “Setting up the environment for Open HyperSwap” on page 101.

Perform these steps to enable HyperSwap or Open HyperSwap for a session.

1. In the IBM Tivoli Storage Productivity Center for Replication navigation tree, select **Sessions**. Click the radio button next to the session that you want to enable.
2. From the **Select Action** menu, select **View/Modify Properties** and click **Go**. If you have not created the session, click **Create Session**. You can enable HyperSwap or Open HyperSwap on the Properties page
3. Under **ESS / DS Metro Mirror Options**, select from the following HyperSwap or Open HyperSwap options:
  - **Manage H1-H2 with HyperSwap**. This option enables a session to manage the H1-H2 sequence using HyperSwap. If you select this option, select from the following additional options.
    - **Disable HyperSwap**. Select this option to prevent a HyperSwap from occurring by command or event.

- **On Configuration Error.** Choose one of the following options:
    - **Partition the system(s) out of the sysplex.** Select this option to partition out of the sysplex when a new system is added to the sysplex and encounters an error in loading the configuration. A restart of your system is required if you select this option.
    - **Disable HyperSwap.** Select this option to prevent a HyperSwap from occurring by command or event.
  - **On Planned HyperSwap Error.** Choose one of the following options:
    - **Partition out the failing system(s) and continue swap processing on the remaining system(s).** Select this option to partition out the failing system and continues the swap processing on any remaining systems.
    - **Disable HyperSwap after attempting backout.** Select this option to stop the HyperSwap action, and disable the HyperSwap commands or events.
  - **On Unplanned HyperSwap.** Choose one of the following options:
    - **Partition out the failing system(s) and continue swap processing on the remaining system(s).** Select this option to partition out the failing systems and continues the HyperSwap processing on the remaining systems when a new system is added to the sysplex and HyperSwap does not complete. A restart of your system is required if you select this option.
    - **Disable HyperSwap after attempting backout.** Select this option to stop the HyperSwap action, and disable the HyperSwap commands or events.
  - **Manage H1-H2 with Open HyperSwap.** If volumes are attached to an IBM AIX host, Tivoli Storage Productivity Center for Replication can manage the H1-H2 sequence of a Metro Mirror session using Open HyperSwap. If this option is selected, a failure on the host accessible volumes triggers a swap, which redirects application I/O to the secondary volumes. Only volumes that are currently attached to the host systems that are defined on the Tivoli Storage Productivity Center for Replication Host Systems panel are eligible for Open HyperSwap.
    - **Disable Open HyperSwap.** Select this option to prevent a swap from occurring by a command or event while keeping the configuration on the host system and all primary and secondary volumes coupled.
4. Click **OK** to apply the selected options.

#### **Restarting an AIX Host System that is enabled for Open HyperSwap:**

When an IBM AIX host system is restarted, the host automatically attempts to open any volumes for input/output (I/O) that were open prior to the restart. If Open HyperSwap was enabled for a set of volumes on the host system, the host must determine which storage system is the primary system before the host can allow the volumes to be opened.

If the Metro Mirror relationship for the set of volumes is in a Prepared or Suspended state and the host has connectivity to both the primary and secondary storage systems, the host can determine through the hardware which storage system is the primary system. In this situation, the host automatically opens the volumes.

If the Metro Mirror relationship for the set of volumes is in a Prepared state and the host has connectivity to only the secondary storage system, all I/O to the

volumes might be blocked on the host system until the host is able to verify the primary volume in the relationship. The AIX command varyonvg will fail to open the volumes for I/O to prevent the application from writing to the incorrect site. If the host can determine which volume is the primary volume in the relationship and connectivity to the primary storage system is still lost, a Hyperswap event is triggered. This event causes all I/O to be automatically opened and directed to the secondary storage system.

If the Metro Mirror relationship for the set of volumes is in a Target Available state after a Hyperswap or a Recover command has been issued for the session, or the host system does not have the connectivity necessary to determine which site is the primary site and all I/O to the volumes are blocked on the host system. The varyonvg command will fail to open the volumes for I/O to prevent the application from writing to the incorrect site.

### **Unblocking I/O on the host system after a host system restart**

When any of the previous scenarios cause I/O to be blocked, manual actions might be necessary to remove the block.

If the relationships are in a Target Available state on the hardware, issue a Start command to the session in the desired direction of the relationship. This action defines the primary storage system for the host. The host system can allow the volumes to be opened to I/O.

If the relationships cannot be restarted, or the host cannot determine the primary storage system, it might be necessary to manually decouple the volumes on the host system.

To decouple the volumes the following options are available:

- Option 1: Terminate the session or remove the copy set. This option would require a full copy when the relationships are restarted.
- Option 2: Remove Object Data Manager (ODM) entries using the following command:

```
odmdelete -o Volume_Equivalency
```

#### **CAUTION:**

**This command should be used only for this scenario because the command deletes copy set information.**

### **Planned and unplanned swaps:**

Once a session has been enabled for HyperSwap or Open HyperSwap and reaches the Prepared state, IBM Tivoli Storage Productivity Center for Replication loads the configuration of volumes that are capable of being swapped onto IBM z/OS or AIX.

When the load is complete, the session is capable of a planned or unplanned swap. The H1-H2 role pair on the session show a type of HS. An H is displayed over the connection in the dynamic image for that role pair as shown in the following

### Metro Mirror Failover/Failback




figure.


### Performing a Planned Swap

Once the session configuration is loaded on z/OS for HyperSwap or AIX for Open HyperSwap, the session is considered swap capable. There may be cases such as a planned maintenance or a migration from the primary storage, in which a planned swap might be required. Once the session is in a swap capable state, a planned swap can be executed by issuing the HyperSwap command against the session.

Once a planned swap is run for z/OS HyperSwap and Open HyperSwap, the session is transitioned to a Target Available state and all the H1-H2 pairs are in a Target Available state. If the H1-H2 role pair was consistent at the time of the swap, the session will have a status of Normal and will indicate that H1-H2 is consistent. If the H1-H2 role pair was not consistent at the time of the swap, the session might display a status of SEVERE because the session is inconsistent. The active host on the session is then displayed as H2.

All input/output (I/O) should have been redirected to the H2 volumes. After a successful swap to site 2, it is not possible to re-enable copy to site 2. Therefore, it is not possible to issue a Start H1->H2 command. The only way to restart the copy is a Start H2->H1 command. To have the volumes protected with high availability and disaster recovery again, the error that caused the swap must be fixed and then the session must be manually restarted to begin copying to the other site.

The following figure illustrates a planned swap.



### What happens when an unplanned swap occurs


Once the session configuration is loaded on z/OS for HyperSwap or AIX for Open HyperSwap, the session is considered swap capable. In the event of a primary I/O error, a swap occurs automatically. For HyperSwap, z/OS performs the entire swap

and then alerts Tivoli Storage Productivity Center for Replication that a swap has occurred. For Open HyperSwap, Tivoli Storage Productivity Center for Replication and the AIX host work together to perform the swap.

Once an unplanned swap occurs for HyperSwap and Open HyperSwap, the session is transitioned to a Target Available state and all the H1-H2 pairs are in a Target Available state. If the H1-H2 role pair was consistent at the time of the swap, the session will have a Status of Normal and will indicate that H1-H2 is consistent. If the H1-H2 role pair was not consistent at the time of the swap, the session might display a status of SEVERE because the session is inconsistent. The active host on the session is then displayed as H2.

All I/O should have been redirected to the H2 volumes. After a successful swap to site 2, it is not possible to re-enable copy to site 2. Therefore, it is not possible to issue a Start H1->H2 command. The only way to restart the copy is a Start H2->H1 command. To have the volumes protected with high availability and disaster recovery again, the error that caused the swap must be fixed and then the session must be manually restarted to begin copying to the other site.

The following figure illustrates an unplanned swap.



### Scenarios requiring a full copy in Metro Global Mirror with HyperSwap sessions

In the following cases, a full copy is required to return to the three-site configuration after a swap:

- If you are running a Metro Global Mirror session with HyperSwap and you issue the HyperSwap command using the z/OS HyperSwap API rather than the Tivoli Storage Productivity Center for Replication graphical user interface (GUI).
- If you are running a Metro Global Mirror session with HyperSwap and an unplanned swap occurs.

*Verifying that a session is capable of a planned or unplanned swap:*

You can verify whether a sessions is capable of a planned or unplanned swap from the IBM z/OS console (HyperSwap) or the IBM AIX host (Open HyperSwap).

Perform these steps to verify the status of HyperSwap from the z/OS console:

1. Issue the `ds hs,status` command for the overall status of the HyperSwap session. For example:

```

15.03.06 SYSTEM1 d hs,status
15.03.06 SYSTEM1 STC00063 IOSHM0303I HyperSwap Status 531
Replication Session: SR_HS
HyperSwap enabled
New member configuration load failed: Disable
Planned swap recovery: Disable
Unplanned swap recovery: Disable
FreezeAll: No
Stop: No

```

2. Issue the `ds hs,config(detail,all)` command to verify all the volumes in the configuration. For example:

```

15.03.51 SYSTEM1 d hs,config(detail,all)
15.03.51 SYSTEM1 STC00063 IOSHM0304I HyperSwap Configuration 534
Replication Session: SR_HS
Prim. SSID UA DEV# VOLSER Sec. SSID UA DEV# Status
06 02 00F42 8K3602 06 04 00FA2
06 01 00F41 8K3601 06 03 00FA1
06 00 00F40 8K3600 06 02 00FA0

```

Perform these steps to verify the status of Open HyperSwap from the AIX host:

1. Issue the `pcmpath query device` command to see the session association and which path the input/output (I/O) is currently being routed to, which is indicated by an asterisk. For example:

```
host1> pcmpath query device 14
```

```

DEV#: 14 DEVICE NAME: hdisk14 TYPE: 2107900 ALGORITHM: Load Balance
SESSION NAME: session1
OS Direction: H1<-H2

```

```
=====
PRIMARY SERIAL: 25252520000

```

| Path# | Adapter/Path Name | State | Mode   | Select | Errors |
|-------|-------------------|-------|--------|--------|--------|
| 0     | fscsi0/path0      | CLOSE | NORMAL | 6091   | 0      |
| 1     | fscsi0/path2      | CLOSE | NORMAL | 6300   | 0      |
| 2     | fscsi1/path4      | CLOSE | NORMAL | 6294   | 0      |
| 3     | fscsi1/path5      | CLOSE | NORMAL | 6187   | 0      |

```
SECONDARY SERIAL: 34343430000 *
```

| Path# | Adapter/Path Name | State | Mode   | Select | Errors |
|-------|-------------------|-------|--------|--------|--------|
| 4     | fscsi0/path1      | CLOSE | NORMAL | 59463  | 0      |
| 5     | fscsi0/path3      | CLOSE | NORMAL | 59250  | 0      |
| 6     | fscsi1/path6      | CLOSE | NORMAL | 59258  | 0      |
| 7     | fscsi1/path7      | CLOSE | NORMAL | 59364  | 0      |

### Temporarily disabling HyperSwap or Open HyperSwap:

In some situations, it might be necessary to temporarily disable the HyperSwap or Open HyperSwap capabilities for a session.

You might want to disable HyperSwap or Open HyperSwap under the following circumstances:

- Performing maintenance
- Inability for one sysplex member to communicate with one or more volumes

Perform these steps to disable Open HyperSwap for a specific session:

1. In the navigation IBM Tivoli Storage Productivity Center for Replication tree, select **Sessions**. The Sessions panel is displayed.
2. Select the sessions for which you want to disable HyperSwap or Open HyperSwap.



3. Select **View/Modify Properties** from the **Select Actions** list, and click **Go**.
4. Select **Disable HyperSwap** or **Disable Open HyperSwap** and click **OK**.

**Tip:** On management servers that run IBM z/OS, you can also disable HyperSwap from an MVS command prompt by entering **SETHS DISABLE**.

### **Using active and standby Tivoli Storage Productivity Center for Replication servers with HyperSwap or Open HyperSwap:**

To ensure that there is an IBM Tivoli Storage Productivity Center for Replication server available in the event of a disaster, an active and standby management server configuration can be set up. You can enable HyperSwap and Open HyperSwap for sessions while maintaining an active and standby server configuration.

#### **Active and standby servers with HyperSwap**

When the storage system is set up to connect through the z/OS interface, the connection information is automatically sent to the standby server and a connection is attempted. It is possible that the connection will fail if the standby server is not running on z/OS or does not have access to the same volumes. If the connection fails, any takeover done on the standby server will not be able to manage the HyperSwap. On z/OS, if the session configuration had been successfully loaded prior to the HyperSwap, the z/OS system is still capable of performing the HyperSwap. If the z/OS system swaps the volumes but there is no communication possible to the Tivoli Storage Productivity Center for Replication server, the session recognizes that the pairs became suspended and the session will go into a Suspended/Severe state. From this state the customer can clear the **Manage H1-H2 with Hyperswap** option and issue the Recover command for the session to get the session into a Target Available state.

#### **Active and standby servers with Open HyperSwap**

When there is an active and standby management server configuration and a host system connection is added to the active server, the host system connection is automatically sent to the standby server and a connection is attempted. Once the session configuration is loaded on IBM AIX, Open HyperSwap is possible only if there is continual communication between AIX and the Tivoli Storage Productivity Center for Replication server. If a takeover is performed on a standby server that is unable to connect to the host system that is managing the swap, the session is no longer Open HyperSwap capable. Communication to the host system must be activated before the session can become Open HyperSwap capable again.

#### **Related tasks:**

“Managing management servers” on page 643

This section provides information about how to set up active and standby management servers, restore a lost connection between the management servers, or perform a takeover on the standby management server.

## **Session commands**

The commands that are available for a session depend on the session type.

Commands are issued synchronously to IBM Tivoli Storage Productivity Center for Replication sessions. Any subsequent command issued to an individual session is not processed until the first command has completed. Some commands can take an extended amount of time to complete, such as the **Start** command as it sets up the

hardware. The GUI continues to allow you to issue commands to other sessions and not hold up functionality. When a command has completed, the console displays the results of the command.

## Basic HyperSwap commands

Use this information to learn about commands available for Basic HyperSwap sessions.

**Note:** Individual suspend and recover commands are not available in HyperSwap.

Table 37. Basic HyperSwap commands

| Command                | Action                                                                                                                                                                                                                                                                                                                                                   |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>HyperSwap</b>       | Triggers a HyperSwap where I/O is redirected from the source volume to the target volume, without affecting the application using those volumes. You can use this command if you want to perform maintenance on the original source volumes.                                                                                                             |
| <b>Start H1-&gt;H2</b> | Starts copying data synchronously from H1 to H2 in a Metro Mirror session.<br><b>Note:</b> A session might go into a Severe state with error code 1000000 before the session returns to Normal/Prepared State and HyperSwap Capable. The duration of the Severe state depends on how large of a session is running.                                      |
| <b>Start H2-&gt;H1</b> | Starts copying data synchronously from H2 to H1 in a Metro Mirror session. You can issue this command only after the session has been swapped and the production site is H2. To enable data protection when the H1 volumes are available again, start I/O to the H2 volumes, and issue this command to replicate data from the H2 volumes to H1 volumes. |
| <b>Stop</b>            | Suspends updates to all the targets of pairs in a session. You can issue this command at any time during an active session.<br><b>Note:</b> After you issue the stop command, targets might not be consistent.                                                                                                                                           |
| <b>Terminate</b>       | Removes all physical copies and relationships from the hardware during an active session.                                                                                                                                                                                                                                                                |

## FlashCopy commands

Use this information to learn about commands available for FlashCopy sessions.

Table 38. FlashCopy commands

| Command      | Action                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Start</b> | Performs any steps necessary to define the relationship before performing a FlashCopy operation. For ESS, DS6000, and DS8000, this command is not an option. Issue this command to put the session in the prepared state for the following storage systems: <ul style="list-style-type: none"> <li>• SAN Volume Controller</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> </ul> |

Table 38. FlashCopy commands (continued)

| Command                       | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Flash</b>                  | <p>Performs the FlashCopy operation using the specified options. Issue the Flash command to create a data consistent point-in-time copy for a FlashCopy Session with volumes on the following storage systems:</p> <ul style="list-style-type: none"> <li>• SAN Volume Controller</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> </ul> <p>For a FlashCopy session containing ESS, DS6000, and DS8000 volumes, the Flash command by itself is not sufficient to create a consistent copy. To create a consistent copy using the ESS, DS6000, and DS8000 Flash commands, you must quiesce application I/O before issuing the Flash command.</p> |
| <b>InitiateBackgroundCopy</b> | <p>Copies all tracks from the source to the target immediately, instead of waiting until the source track is written to. This command is valid <i>only</i> when the background copy is not already running.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Terminate</b>              | <p>Removes all active physical copies and relationships from the hardware during an active session.</p> <p>If you want the targets to be data consistent before removing their relationship, you must issue the <b>InitiateBackgroundCopy</b> command if NOCOPY was specified, and then wait for the background copy to complete by checking the copying status of the pairs.</p>                                                                                                                                                                                                                                                                               |

## Snapshot commands

Use this information to learn about commands that are available for Snapshot sessions and groups. A snapshot group is a grouping of snapshots of individual volumes in a consistency group at a specific point in time.

Table 39. Snapshot session commands

| Command                | Action                                                                                                                                                                                                                                                                                                 |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Create Snapshot</b> | Creates a snapshot of the volumes in the session                                                                                                                                                                                                                                                       |
| <b>Restore</b>         | Restores the H1 volumes in the session from a set of snapshot volumes. You must have at least one snapshot group to restore from. When you issue this command in the Tivoli Storage Productivity Center for Replication graphical user interface (GUI), you are prompted to select the snapshot group. |

Table 40. Snapshot group commands

| Command       | Action                                                                                                                                                                                                                                                              |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Delete</b> | Deletes the snapshot group and all the individual snapshots that are in the group from the session and from the XIV system. If the deleted snapshot group is the last snapshot group that is associated with the session, the session returns to the Defined state. |

Table 40. Snapshot group commands (continued)

| Command             | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Disband</b>      | Disbands the snapshot group. When a snapshot group is disbanded, the snapshot group no longer exists. All snapshots in the snapshot group become individual snapshots that are no longer associated to the consistency group or the session. After a snapshot group is disbanded, it is no longer displayed in or managed by Tivoli Storage Productivity Center for Replication. If the disbanded snapshot group is the last snapshot group that is associated with the session, the session returns to the Defined state. |
| <b>Duplicate</b>    | Duplicates the snapshot group. When a snapshot group is duplicated, a new snapshot group is created with new snapshots for all volumes that are in the duplicated group. The name of the duplicated snapshot group is generated automatically by the XIV system.                                                                                                                                                                                                                                                           |
| <b>Lock</b>         | Locks a snapshot group. If the snapshot group is locked, write operations to the snapshots that are in the snapshot group are prevented. By default, a snapshot group is locked when it is created. This action is valid only if the snapshot group is unlocked.                                                                                                                                                                                                                                                           |
| <b>Overwrite</b>    | Overwrites the snapshot group to reflect the data that is on the H1 volume.                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Rename</b>       | Renames the snapshot group to a name that you provide. The name can be a maximum of 64 alphanumeric characters.                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Restore</b>      | Restores the contents of a snapshot group by using another snapshot group in the session. Both of the snapshot groups must contain the same subset of volumes.                                                                                                                                                                                                                                                                                                                                                             |
| <b>Set Priority</b> | Sets the priority in which a snapshot group is deleted. The value can be the number 1 - 4. A value of 1 specifies that the snapshot group is deleted last. A value of 4 specifies that the snapshot group is deleted first.                                                                                                                                                                                                                                                                                                |
| <b>Unlock</b>       | Unlocks a snapshot group. If the snapshot group is unlocked, write operations to the snapshots that are in the snapshot group are enabled and the snapshot group is displayed as modified. This action is valid only if the snapshot group is locked.                                                                                                                                                                                                                                                                      |

## Metro Mirror commands

Use this information to learn about commands available for Metro Mirror sessions.

Table 41. Metro Mirror commands

| Command                      | Action                                                                                                                                                                                                                                                 |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Enable Copy to Site 1</b> | Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the <b>Start H2-&gt;H1</b> command becomes available. |
| <b>Enable Copy to Site 2</b> | Run this command and confirm that you want to reverse the direction of replication before reversing the direction of copying in a failover and failback session. After you issue this command, the <b>Start H1-&gt;H2</b> command becomes available.   |

Table 41. Metro Mirror commands (continued)

| Command                | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>HyperSwap</b>       | Triggers a HyperSwap where I/O is redirected from the source volume to the target volume, without affecting the application using those volumes. You can use this command if you want to perform maintenance on the original source volumes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Start</b>           | Establishes a single-direction session with the hardware and begins the synchronization process between the source and target volumes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Start H1-&gt;H2</b> | Establishes Metro Mirror relationships between the H1 volumes and the H2 volumes, and begins data replication from H1 to H2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Start H2-&gt;H1</b> | Establishes Metro Mirror relationships between the H2 volumes and the H1 volumes and starts data replication from H2 to H1. Indicates direction of a failover and failback between two hosts in a Metro Mirror session. If the session has been recovered such that the production site is now H2, you can issue the <b>Start H2-&gt;H1</b> command to start production on H2 and start data replication.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Stop</b>            | Inconsistently suspends updates to all the targets of pairs in a session. This command can be issued at any point during an active session.<br><b>Note:</b> Targets after the suspend are not considered to be consistent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>StartGC</b>         | Establishes Global Copy relationships between the H1 volumes and the H2 volumes, and begins asynchronous data replication from H1 to H2. While in the Preparing state, it will not change to the Prepared state unless you switch to Metro Mirror.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Suspend</b>         | Causes all target volumes to remain at a data-consistent point and stops all data that is moving to the target volumes. This command can be issued at any point during a session when the data is actively being copied.<br><b>Note:</b> It is recommended that you avoid using the same LSS pairs for multiple Metro Mirror sessions. Metro Mirror uses a freeze command on ESS, DS6000, and DS8000 storage systems to create the data-consistent point. If there are other Metro Mirror sessions overlapping the same LSS pairs as in this session, those sessions are also suspended.<br><br>When a <b>Suspend</b> command is issued to a source volume in an LSS that has source volumes in another active Metro Mirror session, the other source volumes are affected only if they have the same target LSS. The primary volumes are suspended, but volumes in the same source LSS that have target volumes in a different LSS are not affected because they use a different PPRC path connection. |
| <b>Recover</b>         | Issues the <b>Recover</b> command to suspended sessions. This command performs the steps necessary to make the target available as the new primary site. Upon completion of this command, the session becomes Target Available.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Terminate</b>       | Removes all copy relationships from the hardware during an active session. If you want the targets to be data consistent before removing their relationship, you must issue the <b>Suspend</b> command, then the <b>Recover</b> command, and then the <b>Terminate</b> command.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## Metro Mirror with Practice commands

Use this information to learn about commands available for Metro Mirror with Practice sessions.

Table 42. Metro Mirror with Practice commands

| Command               | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enable Copy to Site 1 | Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the <b>Start H2-&gt;H1</b> command becomes available.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Enable Copy to Site 2 | Run this command and confirm that you want to reverse the direction of replication before reversing the direction of copying in a failover and failback session. After you issue this command, the <b>Start H1-&gt;H2</b> command becomes available.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Flash                 | Creates a FlashCopy image from I2 volumes to H2 volumes. The amount of time for this to occur will vary depending on the number of copy sets in the session.<br><b>Note:</b> For ESS, DS6000, DS8000 storage systems, the <b>Flash</b> command uses the freeze and thaw processing to create a data consistent point for the FlashCopy. If there is another Metro Mirror session overlapping on one or more of the same LSS pairs, that session will be suspended. It is also possible that the suspension of the other session might cause the Metro Mirror session to remain suspended after the flash command is issued instead of returning to Prepared state. Avoid using the same LSS pairs for multiple Metro Mirror sessions if possible. |
| Start H1->H2          | Establishes a Metro Mirror relationship from the H1 volumes to the I2 volumes, and begins data replication.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Start H2->H1          | Establishes a Metro Mirror relationship from H2 to H1 and begins data replication.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| StartGC_H1H2          | Distinguishes when the session is in the Preparing state from H1 to I2 and begins the asynchronous process between the source and target volumes. While in the Preparing state the session will not change to the Prepared state unless you switch to Metro Mirror.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| StartGC_H2H1          | Distinguishes when the session is in the Preparing state from H2 to H1 and begins the asynchronous process between the source and target volumes. While in the Preparing state the session will not change to the Prepared state unless you switch to Metro Mirror.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

Table 42. Metro Mirror with Practice commands (continued)

| Command          | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Suspend</b>   | <p>Causes all target volumes to remain at a data-consistent point and stops all data that is moving to the target volumes. This command can be issued at any point during a session when the data is actively being copied.</p> <p><b>Note:</b> The Metro Mirror command uses a freeze command on the ESS, DS6000, or DS8000 devices to create the data-consistent point. If there are other Metro Mirror sessions overlapping the same LSS pairs as in this session, those sessions will also become suspended. Avoid using the same LSS pairs for multiple Metro Mirror sessions.</p> <p>When a <b>Suspend</b> command is issued to a source volume in an LSS that has source volumes in another active Metro Mirror session, the other source volumes are affected only if they have the same target LSS. The primary volumes are suspended, but volumes in the same source LSS that have target volumes in a different LSS are not affected because they use a different PPRC path connection.</p> |
| <b>Stop</b>      | <p>Inconsistently suspends updates to all the targets of pairs in a session. This command can be issued at any point during an active session.</p> <p><b>Note:</b> Targets after the suspend are not considered to be consistent.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Terminate</b> | Terminates all copy relationships on the hardware.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Recover</b>   | <p>Takes a point-in-time copy of the data on I2 to the H2 volumes, enabling the application to be attached and run from the H2 volumes on site 2.</p> <p><b>Note:</b> The point-in-time copy is performed when the session is in a recoverable state to avoid that previous consistent data on H2 are overwritten by inconsistent data upon Recover. You can issue the <b>Flash</b> command if you want to force a point-in-time copy from I2 to JH2 volumes afterwards.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## Global Mirror commands

Use this information to learn about commands available for Global Mirror sessions.

Table 43. Global Mirror commands

| Command                      | Action                                                                                                                                                                                                                                                 |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Enable Copy to Site 1</b> | Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the <b>Start H2-&gt;H1</b> command becomes available. |
| <b>Enable Copy to Site 2</b> | Run this command and confirm that you want to reverse the direction of replication before reversing the direction of copying in a failover and failback session. After you issue this command, the <b>Start H1-&gt;H2</b> command becomes available.   |
| <b>Start</b>                 | Establishes all relationships in a single-direction session and begins the process necessary to start forming consistency groups on the hardware.                                                                                                      |

Table 43. Global Mirror commands (continued)

| Command                  | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Start H1-&gt;H2</b>   | Starts copying data from H1 to H2 in a Global Mirror failover and failback session. Establishes the necessary relationships in the session and begins the process necessary to start copying data from the H1 site to the H2 site and to start forming consistency groups.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Start H2-&gt;H1</b>   | Starts copying data from H2 to H1 in a failover and failback session for ESS, DS6000 and DS8000 sessions. If a recover has been performed on a session such that the production site is now H2, you can issue a <b>Start H2-&gt;H1</b> to start moving data back to Site 1. However, this start does not provide consistent protection as it copies only asynchronously back because of the long distance. A Global Copy relationship is used. When you are ready to move production back to Site 1, issue a suspend to the session; this puts the relationships into a synchronized state and suspends them consistently. Sessions are consistent when copying H2->H1 for the following storage systems: <ul style="list-style-type: none"> <li>• SAN Volume Controller</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> </ul> |
| <b>StartGC H1-&gt;H2</b> | Establishes Global Copy relationships between site 1 and site 2 and begins asynchronous data replication from H1 to H2. To change the session state from Preparing to Prepared, you must issue the Start H1->H2 command and the session must begin to form consistency groups. <p>There is no disaster recovery protection for Global Copy relationships. If a disaster such as the loss of a primary storage system or a link failure between the sites occurs, the session might be inconsistent when you issue the Recover command.</p> <p>This command is available for Global Mirror Failover/Failback sessions for the following storage systems:</p> <ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS6000</li> <li>• System Storage DS8000</li> </ul>              |
| <b>Suspend</b>           | Stops all consistency group formation when the data is actively being copied and suspends the H1->H2 Global Copy pairs. To issue the pause command to the Global Mirror session on the hardware without suspending the Global Copy pairs, open the <code>rmserver.properties</code> file and add the following property to disable the Global Copy suspension on the <b>Suspend</b> command: <b>csm.server.sus_gc_pairs_on_gm_pause = false</b> . The default property is <b>true</b> and the <b>Suspend</b> command automatically suspends the Global Copy pairs.                                                                                                                                                                                                                                                                              |
| <b>Recover</b>           | Issue this command to recover the session to the target site. This command performs the steps necessary to make the target host volumes consistent and available for access as the new primary site. Upon completion of this command, the session becomes Target Available. Do not access H2 volumes until the <b>Recover</b> command is completed and the session displays Target Available and Recoverable. A Recover to H2 also establishes a point-in-time copy to J2 to preserve the last consistent data.                                                                                                                                                                                                                                                                                                                                 |



Table 43. Global Mirror commands (continued)

| Command          | Action                                                                                                                                                                                                                                                                                           |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Terminate</b> | Removes all physical copies and relationships from the hardware during an active session.<br><br>If you want the targets to be data consistent before removing their relationship, you must issue the <b>Suspend</b> command, the <b>Recover</b> command, and then the <b>Terminate</b> command. |

## Global Mirror with Practice commands

Use this information to learn about commands available for Global Mirror with Practice sessions.

Table 44. Global Mirror with Practice commands

| Command                      | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Enable Copy to Site 1</b> | Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the <b>Start H2-&gt;H1</b> command becomes available.                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Enable Copy to Site 2</b> | Run this command and confirm that you want to reverse the direction of replication before reversing the direction of copying in a failover and failback session. After you issue this command, the <b>Start H1-&gt;H2</b> command becomes available.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Flash</b>                 | The Flash command on a Global Mirror Practice session for ESS, DS6000, and DS8000 temporarily pauses the formation of consistency groups, ensure that all I2s are consistent, and then flash the data from I2 to the H2 volumes. After the flash is complete, the Global Mirror session is automatically restarted, and the session begins forming consistency groups on I2. You can then use the H2 volumes to practice your disaster recovery procedures.                                                                                                                                                                                                                                       |
| <b>Start H1-&gt;H2</b>       | Starts copying data from H1 to H2. After the first pass of the copy is complete for all pairs, the session establishes the I2->J2 FlashCopy pairs, and starts the Global Mirror master so that the hardware will begin forming consistency groups, to ensure consistent data is at site 2.                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Start H2-&gt;H1</b>       | Starts copying data from H2 to H1 in a failover and failback session. If a recover has been performed on a session such that the production site is now H2, you can issue a <b>Start H2-&gt;H1</b> to start moving data back to Site 1. However, this start does not provide consistent protection as it copies only asynchronously back because of the long distance.<br><b>Note:</b> ESS, DS6000, and DS8000 volumes are not consistent for the <b>Start H2-&gt;H1</b> command. A Global Copy relationship is used. When you are ready to move production back to Site 1, issue a suspend to the session; this puts the relationships into a synchronized state and suspends them consistently. |

Table 44. Global Mirror with Practice commands (continued)

| Command                  | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>StartGC H1-&gt;H2</b> | <p>Establishes Global Copy relationships between site 1 and site 2 and begins asynchronous data replication from H1 to I2. To change the session state from Preparing to Prepared, you must issue the Start H1-&gt;H2 command and the session must begin to form consistency groups.</p> <p>There is no disaster recovery protection for Global Copy relationships. If a disaster such as the loss of a the primary Tivoli Storage Productivity Center for Replication server occurs, the session might be inconsistent when you issue the Recover command.</p> <p>This command is available for Global Mirror Failover/Failback and Global Mirror Failover/Failback with Practice sessions for the following storage systems:</p> <ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server Model 800</li> <li>• System Storage DS6000</li> <li>• System Storage DS8000</li> </ul> |
| <b>Terminate</b>         | Removes all physical copies and relationships on the hardware.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Suspend</b>           | <p>Stops all consistency group formation when the data is actively being copied and suspends the H1-I2 Global Copy pairs. To issue the pause command to the Global Mirror session on the hardware without suspending the Global Copy pairs, open the <code>rmserver.properties</code> file and add the following property to disable the Global Copy suspension on the <b>Suspend</b> command: <b><code>csm.server.sus_gc_pairs_on_gm_pause = false</code></b>. The default property is <b>true</b> and the <b>Suspend</b> command automatically suspends the Global Copy pairs..</p>                                                                                                                                                                                                                                                                                                                     |
| <b>Recover</b>           | <p>Restores consistent data on I2 volumes and takes a point-in-time copy of the data on I2 to the H2 volumes, enabling the application to be attached and run from the H2 volumes on site 2. The I2 Volumes will continue to hold the consistent data and can be flashed again to H2 using the <b>Flash</b> command.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## Metro Global Mirror commands

Use this information to learn about commands available for Metro Global Mirror sessions.

Table 45. Metro Global Mirror commands

| Command                      | Action                                                                                                                                                                                                                                                        |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Enable Copy to Site 1</b> | Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. After you issue this command, the <b>Start H2-&gt;H1-&gt;H3</b> command becomes available. |
| <b>Enable Copy to Site 2</b> | Run this command and confirm that you want to reverse the direction of replication before reversing the direction of copying in a failover and failback session. After you issue this command, the <b>Start H1-&gt;H2-&gt;H3</b> command becomes available.   |

Table 45. Metro Global Mirror commands (continued)

| Command                       | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>HyperSwap</b>              | <p>Causes a site switch, equivalent to a suspend and recover for a Metro Mirror with failover and failback individual suspend and recover commands are not available.</p> <p>The Metro Global Mirror for HyperSwap supports IBM Tivoli Storage Productivity Center for Replication installed on z/OS with the IBM Tivoli Storage Productivity Center for Replication Three Site Business Continuity license.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Start H1-&gt;H2-&gt;H3</b> | <p>(This is the Metro Global Mirror initial start command.) Establishes Metro Mirror relationships between H1 and H2, and Global Mirror relationships between H2 and H3. For Metro Global Mirror, this includes the J3 volume to complete the Global Mirror configuration. (The J3 volume role is the journal volume at site 3.). <b>Start H1-&gt;H2-&gt;H3</b> can be used from some Metro Global Mirror configurations to transition back to the starting H1-&gt;H2-&gt;H3 configuration.</p> <p>This command is valid only when the session is in a defined, preparing, prepared, or suspended state.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Start H1-&gt;H3</b>        | <p>From the H1-&gt;H2-&gt;H3 configuration, this command changes the session configuration to a Global Mirror-only session between H1 and H3, with H1 as the source. Use this command in case of an H2 failure with transition bitmap support provided by incremental resynchronization. It can be used when a session is in preparing, prepared, and suspended states because there is not a source host change involved.</p> <p>This command allows you to bypass the H2 volume in case of an H2 failure and copy only the changed tracks and tracks in flight from H1 to H3. After the incremental resynchronization is performed, the session is running Global Mirror from H1 to H3 and thus loses the near-zero data loss protection achieved with Metro Mirror when running H1-&gt;H2-&gt;H3. However, data consistency is still maintained at the remote site with the Global Mirror solution.</p> <p>From H2-&gt;H1-&gt;H3 configuration, this command changes the session configuration to a Global Mirror-only session configuration between H1 and H3, with H1 as the source. Use this command when the source site has a failure and production is moved to the H1 site. This can be done for unplanned HyperSwap. The Global Mirror session is continued. This is a host-volume change so this command is valid only when restarting the H1-&gt;H3 configuration or from the TargetAvailable H2-&gt;H1-&gt;H3 state.</p> |

Table 45. Metro Global Mirror commands (continued)

| Command                       | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Start H2-&gt;H3</b>        | <p>From the H1-&gt;H2-&gt;H3 configuration, this command moves the session configuration to a configuration between H2 and H3, with H2 as the source. Use this command when the source site has a failure and production is moved to the H2 site, for example, for unplanned HyperSwap. The Global Mirror session is continued. This session is a host-volume change so this command is valid only when restarting the H1-&gt;H3 configuration or from the TargetAvailable H2-&gt;H1&gt;H3 state.</p> <p>From the H2-&gt;H1-&gt;H3 configuration, this command changes the session configuration to a configuration between H2 and H3 with H2 as the source. Use this command in case of an H1 failure with transition bitmap support provided by incremental resynchronization. Because there is not a source-host change involved, it can be used when the session is in preparing, prepared, and suspended states <b>Start H2-&gt;H1-&gt;H3</b> can be used to transition back to the starting H2-&gt;H1-&gt;H3 configuration.</p> |
| <b>Start H2-&gt;H1-&gt;H3</b> | <p>(This is the Metro Global Mirror start command.) This is the configuration that completes the HyperSwap processing. This command creates Metro Mirror relationships between H2 and H1 and Global Mirror relationships between H1 and H3. For Metro Global Mirror, this includes the J3 volume to complete the Global Mirror configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Start H3-&gt;H1-&gt;H2</b> | <p>After recovering to H3, this command sets up the hardware to allow the application to begin writing to H3, and the data is copied back to H1 and H2. However, issuing this command does not guarantee consistency in the case of a disaster because only Global Copy relationships are established to cover the long distance copy back to site 1.</p> <p>To move the application back to H1, you can issue a suspend while in this state to drive all the relationships to a consistent state and then issue a freeze to make the session consistent. You can then issue a <b>Rcover</b> followed by a <b>Start H1-&gt;H2-&gt;H3</b> to go back to the original configuration.</p>                                                                                                                                                                                                                                                                                                                                                |
| <b>SuspendH2H3</b>            | <p>When running H1-&gt;H2-&gt;H3, this command issues a pause to the Global Mirror master and causes the Global Mirror master to stop forming consistency groups.</p> <p>This command is valid only when the session is in a prepared state. Stops all consistency group formation when the data is actively being copied and suspends the H2&gt;H3 Global Copy pairs. To issue the pause command to the Global Mirror session on the hardware without suspending the Global Copy pairs, open the <code>rmserver.properties</code> file and add the following property to disable the Global Copy suspension on the <b>Suspend</b> command: <b>csm.server.sus_gc_pairs_on_gm_pause = false</b>. The default property is <b>true</b> and the <b>Suspend</b> command automatically suspends the Global Copy pairs..</p>                                                                                                                                                                                                                 |

Table 45. Metro Global Mirror commands (continued)

| Command            | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>SuspendH1H3</b> | <p>When running H2-&gt;H1-&gt;H3, this command issues a pause to the Global Mirror master and causes the Global Mirror master to stop forming consistency groups.</p> <p>This command is valid only when the session is in a prepared state. Stops all consistency group formation when the data is actively being copied and suspends the H2-&gt;H3 Global Copy pairs. To issue the pause command to the Global Mirror session on the hardware without suspending the Global Copy pairs, open the <code>rmsserver.properties</code> file and add the following property to disable the Global Copy suspension on the <b>Suspend</b> command: <b><code>csm.server.sus_gc_pairs_on_gm_pause = false</code></b>. The default property is <b>true</b> and the <b>Suspend</b> command automatically suspends the Global Copy pairs..</p> |
| <b>RecoverH1</b>   | <p>Specifying H1 makes the H1 volume TargetAvailable. Metro Global Mirror (when running H2-&gt;H1-&gt;H3) can move production to either the H1 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing differs, depending on the recovery site. Therefore, the site designation is added to the <b>Recover</b> command so IBM Tivoli Storage Productivity Center for Replication can set up for the failback.</p>                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>RecoverH2</b>   | <p>Specifying H2 makes the H2 volume TargetAvailable. Metro Global Mirror (when running H1-&gt;H2-&gt;H3) can move production to either the H2 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing is different depending on the recovery site. Therefore the site designation is added to the <b>Recover</b> command so IBM Tivoli Storage Productivity Center for Replication can prepare for the failback.</p>                                                                                                                                                                                                                                                                                                                                                                                |
| <b>RecoverH3</b>   | <p>Specifying H3 makes the H3 volume TargetAvailable. Metro Global Mirror (when running H1-&gt;H2-&gt;H3) can then move production to the H3 set of volumes. Because IBM Tivoli Storage Productivity Center for Replication processing differs depending on the recovery site, the site designation is added to the <b>Recover</b> command so that IBM Tivoli Storage Productivity Center for Replication can prepare for the failback.</p> <p>This command prepares H3 so that you can start the application on H3. H3 becomes the active host, and you then have the option start H3-&gt;H1-&gt;H2 to perform a Global Copy copy back. The recovery establishes point-in-time copy to J3 volumes to preserve the last consistent data.</p>                                                                                         |

## Metro Global Mirror with Practice commands

Use this information to learn about commands available for Metro Global Mirror with Practice sessions.

Table 46. Metro Global Mirror with Practice commands

| Command                      | Action                                                                                                                                                             |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Enable Copy to Site 1</b> | Run this command and confirm that you want to reverse the direction of replication before you reverse the direction of copying in a failover and failback session. |
| <b>Enable Copy to Site 2</b> | Run this command and confirm that you want to reverse the direction of replication before reversing the direction of copying in a failover and failback session.   |

Table 46. Metro Global Mirror with Practice commands (continued)

| Command                         | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Flash</b>                    | <p>This command is available in the following states:</p> <ul style="list-style-type: none"> <li>• Target Available state when the active host is H3<br/> <b>Note:</b> Use this command if the FlashCopy portion of the Recover command from I3 to H3, fails for any reason. The problem can be addressed; and a <b>Flash</b> command issued to complete the flash of the consistent data from I3 to H3.</li> <li>• Prepared state when the active host is H1 and data is copying H1 to H2 to I3, or the active host is H2 and data is copying H2 to H1 to H3.</li> <li>• Prepared state when the active host is H2 and data is copying H2 to I3.</li> <li>• Prepared state when the active host is H1 and data is copying H1 to I3.</li> </ul> <p>Use this command if the FlashCopy portion of the Recover command from I3 to H3, fails for any reason. The problem can be addressed, and a <b>Flash</b> command issued to complete the flash of the consistent data from I3 to H3.</p> <p>Issuing a <b>Flash</b> command on a Global Mirror Practice session for ESS, DS6000, and DS8000 will temporarily pause the formation of consistency groups, ensure that all I3s are consistent, and then flash the data from I3 to the H3 volumes. After the flash is complete, the Global Mirror session will be automatically restarted, and the session will begin forming consistency groups on I3. You can then use the H3 volumes to practice your disaster recovery procedures.</p> |
| <b>RecoverH1</b>                | <p>Specifying H1 makes the H1 volume TargetAvailable. Metro Global Mirror (when running H1-&gt;H2-&gt;H3) can move production to either the H2 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing differs, depending on the recovery site. Therefore the site designation is added to the <b>Recover</b> command so IBM Tivoli Storage Productivity Center for Replication can prepare for the failback. The FlashCopy creates a consistent copy of the data on the H3 volumes so that an application can recover to those volumes and begin writing I/O. When the FlashCopy is complete, the session will reach a Target Available state, and you can attach your volumes on Site 3.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>RecoverH2</b>                | <p>Specifying H2 makes the H2 volume TargetAvailable. Metro Global Mirror (when running H1-&gt;H2-&gt;H3) can move production to either the H2 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing differs, depending on the recovery site. Therefore the site designation is added to the <b>Recover</b> command so IBM Tivoli Storage Productivity Center for Replication can prepare for the failback.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>RecoverH3</b>                | <p>Specifying H3 makes the H3 volume the TargetAvailable. When running H1-&gt;H2-&gt;H3, Metro Global Mirror can move production to either the H2 or H3 set of volumes. IBM Tivoli Storage Productivity Center for Replication processing differs, depending on the recovery site; therefore, the site designation is added to the <b>Recover</b> command so IBM Tivoli Storage Productivity Center for Replication can prepare for the failback.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Re-enable Copy to Site 1</b> | <p>After issuing a <b>RecoverH1</b> command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

Table 46. Metro Global Mirror with Practice commands (continued)

| Command                         | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Re-enable Copy to Site 2</b> | After issuing a <b>RecoverH2</b> command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Re-enable Copy to Site 3</b> | After issuing a <b>RecoverH3</b> command, you can run this command to restart the copy to the original the direction of replication in a failover and failback session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Start H1-&gt;H2-&gt;H3</b>   | <p>Metro Global Mirror initial start command. This command creates Metro Mirror relationships between H1 and H2, and Global Mirror relationships between H2 and H3. For Metro Global Mirror, this includes the J3 volume to complete the Global Mirror configuration. (The J3 volume role is the journal volume at site 3.). <b>Start H1-&gt;H2-&gt;H3</b> can be used from some Metro Global Mirror configurations to return to the starting H1&gt;H2&gt;H3 configuration.</p> <p>This command is valid only when the session is in a defined, preparing, prepared, target-available, or suspended state.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Start H1-&gt;H3</b>          | <p>From the H1-&gt;H2-&gt;H3 configuration, this command changes the session configuration to a Global-Mirror-only session between H1 and H3, with H1 as the source. Use this command in case of an H2 failure with transition bitmap support provided by incremental resynchronization. Because there is not a source host change involved, it can be used when a session is in preparing, prepared, and suspended states.</p> <p>You can use this command to bypass the H2 volume in case of an H2 failure and copy only the changed tracks and tracks in flight from H1 to H3. After the incremental resynchronization is performed, the session is running Global Mirror from H1 to H3 and thus loses the near-zero data loss protection achieved with Metro Mirror when running H1-&gt;H2-&gt;H3. However, data consistency is still maintained at the remote site with the Global Mirror solution.</p> <p>From H2-&gt;H1-&gt;H3 configuration, this command changes the session configuration to a Global-Mirror-only session configuration between H1 and H3, with H1 as the source. Use this command when the source site has a failure and production is moved to the H1 site. This can be done for unplanned HyperSwap. The Global Mirror session is continued. This is a host-volume change so this command is valid only when restarting the H1-&gt;H3 configuration or from the TargetAvailable H2-&gt;H1-&gt;H3 state.</p> |

Table 46. Metro Global Mirror with Practice commands (continued)

| Command                | Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Start H2-&gt;H3</b> | <p>From the H1-&gt;H2-&gt;H3 configuration, this command moves the session configuration to a configuration between H2 and H3, with H2 as the source. Use this command when the source site has a failure and production is moved to the H2 site. This can be done for unplanned HyperSwap. The Global Mirror session is continued. This is a host-volume change so this command is valid only when restarting the H1&gt;H3 configuration or from the TargetAvailable H2-&gt;H1-&gt;H3 state.</p> <p>From the H2-&gt;H1-&gt;H3 configuration, this command changes the session configuration to a configuration between H2 and H3 with H2 as the source. Use this command in case of an H1 failure with transition bitmap support provided by incremental resynchronization. Because there is not a source-host change involved it can be used when the session is in preparing, prepared, and suspended states. <b>Start H2-&gt;H1-&gt;H3</b> can be used to return to the starting H2-&gt;H1-&gt;H3 configuration.</p> |
| <b>SuspendH2H3</b>     | <p>When running H1-&gt;H2-&gt;H3, this command issues a pause to the Global Mirror master and causes the Global Mirror master to stop forming consistency groups.</p> <p>This command is valid only when the session is in a prepared state. It stops all consistency group formation when the data is actively being copied and suspends the H2-&gt;H3 Global Copy pairs. To issue the pause command to the Global Mirror session on the hardware without suspending the Global Copy pairs, open the <code>rmsvr.properties</code> file and add the following property to disable the Global Copy suspension on the <b>Suspend</b> command: <b><code>csm.server.sus_gc_pairs_on_gm_pause = false</code></b>. The default property is <b>true</b> and the <b>Suspend</b> command automatically suspends the Global Copy pairs..</p>                                                                                                                                                                                      |
| <b>SuspendH1H3</b>     | <p>When running H2-&gt;H1-&gt;H3, this command issues a pause to the Global Mirror master and causes the Global Mirror master to stop forming consistency groups.</p> <p>This command is valid only when the session is in a prepared state. Stops all consistency group formation when the data is actively being copied and suspends the H2-&gt;H3 Global Copy pairs. To issue the pause command to the Global Mirror session on the hardware without suspending the Global Copy pairs, open the <code>rmsvr.properties</code> file and add the following property to disable the Global Copy suspension on the <b>Suspend</b> command: <b><code>csm.server.sus_gc_pairs_on_gm_pause = false</code></b>. The default property is <b>true</b> and the <b>Suspend</b> command automatically suspends the Global Copy pairs..</p>                                                                                                                                                                                         |
| <b>Terminate</b>       | This command terminates all copy relationships on the hardware.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## Site awareness

You can associate a location with each storage system and each site in a session. This *site awareness* ensures that only the volumes whose location matches the location of the site are allowed for selection when you add copy sets to the session. This prevents a session relationship from being established in the wrong direction.

**Note:** To filter the locations for site awareness, you must first assign a site location to each storage system.



IBM Tivoli Storage Productivity Center for Replication does not perform automatic discovery of locations. Locations are user-defined and specified manually.

You can change the location associated with a storage system that has been added to the IBM Tivoli Storage Productivity Center for Replication configuration. You can choose an existing location or add a new one. Locations are deleted when there is no longer a storage system with an association to that location.

When adding a copy set to a session, a list of candidate storage systems is presented, organized by location. Storage systems that do not have a location are displayed and available for use when you create a copy set.

You can also change the location for any site in a session. Changing the location of a session does not affect the location of the storage systems that are in the session.

Changing the location of a storage system might have consequences. When a session has a volume role with a location that is linked to the location of the storage system, changing the location of the storage system could change the session's volume role location. For example, if there is one storage system with the location of A\_Location and a session with the location of A\_Location for its H1 role, changing the location of the storage system to a different location, such as B\_Location, also changes the session's H1 location to Site 1. However, if there is a second storage system that has the location of A\_Location, the session's role location is not changed.

**Important:** Location matching is enabled only when adding copy sets. If you change the location of a storage system or volume role, IBM Tivoli Storage Productivity Center for Replication does not audit existing copy sets to confirm or deny location mismatches.

## Preserve Mirror option

This topic presents recommendations for using the Preserve Mirror option in FlashCopy and Metro Mirror sessions.

When the source of the FlashCopy relationship is a source of a Metro Mirror relationship, and the target of the FlashCopy relationship is the source of a Metro Mirror relationship, the Preserve Mirror option attempts to preserve the consistency of the Metro Mirror relationship at the target of the FlashCopy relationship, preventing a full copy from being performed over the Metro Mirror link. Instead, parallel flashes will be performed (if possible) on both sites. If the consistency cannot be preserved, the Flash for the FlashCopy relationships will fail, and the data of the Metro Mirror relationship at the target of the FlashCopy relationship will not be changed.

**Note:** This option is available only on DS8000 storage devices with the required code levels installed.

However, in some instances, the Preserve Mirror option can cause a Metro Mirror session to go into a Preparing state, or even a Suspended state. This topic describes the recommended usage of the Preserve Mirror feature. Using this feature in other ways might lead to a Metro Mirror session going into a Preparing or Suspended state.

## FlashCopy session

You can use the Preserve Mirror option in FlashCopy sessions in two different methods:

### Perform an incremental resynchronization

To perform an incremental resynchronization, select the **Incremental** and **Persistent** options in the FlashCopy session: do *not* select the **No Copy** option.

### Perform a single full copy

To perform a single full copy, ensure that the **Incremental**, **Persistent** and **No Copy** options are *not* selected before you issue a **Flash** command. If you use the **No Copy** option, issue either an **Initiate Background Copy** command or **Terminate** command before you issue the Flash command.

Refer to your DS8000 documentation for more information about the Preserve Mirror function.

## Metro Mirror session

You can set up your Metro Mirror pairs in two different ways, depending on the level of consistency you need, and your preferences.

**Note:** For the examples in this section, the source pair is H1a->H2a and the target pair is H1b->H2b. The source pair will contain volumes that will be the source of the FlashCopy relationship and the target pair will always contain volumes that will be the target of the FlashCopy relationship.

### Create one Metro Mirror session, and add the Metro Mirror pairs as copy sets to that session

The benefit to this approach is that you do not need to worry about whether the host considers the H1a->H2a and H1b->H2b volumes to be consistent with one another. IBM Tivoli Storage Productivity Center for Replication will ensure that all of the volumes remain consistent.

A drawback to this approach is that when using the **Attempt to preserve Metro Mirror consistency, but fail FlashCopy if Metro Mirror target consistency cannot be preserved** option (Preserve Mirror Required), there is a chance that the target pair (H1b->H2b) might suspend unexpectedly: this causes all other pairs in the Metro Mirror session to suspend (including H1a->H2a). This can occur when a FlashCopy establish or withdraw fails unexpectedly on the remote (H1b->H2b) site. If the host requires the H1a->H2a and H1b->H2b volumes to be consistent, then you should suspend all other volumes.

### Create one Metro Mirror session for the H1a->H2a volumes, and another Metro Mirror session for the H1b->H2b volumes

Use this option when the hosts and applications do not require the H1a->H2a and H1b->H2b volumes to be consistent with one another. In this case, you should create one Metro Mirror session for all of the H1a->H2a volumes, and another Metro Mirror session for the H1b->H2b volumes. The H1a->H2a pair is added to the first session, while the H1b->H2b pair is added to the second Metro Mirror session. As long as the host does not require consistency between the H1a and H1b volumes, this option benefits you when you use the **Attempt to preserve Metro Mirror consistency, but fail FlashCopy if Metro Mirror target consistency cannot be preserved** option (Preserve Mirror Required). The benefit is that if one

pair is suspended (such as H1a->H2a), the pairs in the other session will not be affected, since it is in a different Metro Mirror session. Using this method, you can avoid the situation in which a critical application is writing to the source pair (H1a->H2a), while a batch job is writing to the target pair (H1b->H2b), and both pairs are in the same IBM Tivoli Storage Productivity Center for Replication session. These factors cause both applications to receive extended long busy signals, instead of just the batch job.

## Creating sessions and adding copy sets

This section describes how to create a session for a specific replication method and then add copy sets to that session.

### Creating a FlashCopy session and adding copy sets

FlashCopy replication creates a point-in-time copy in which the target volume contains a copy of the data that was on the source volume when the FlashCopy session was established.

When you create a FlashCopy session for Global Mirror or Metro Global Mirror space-efficient target volumes, you must select the **No Copy** option for the FlashCopy session. With space-efficient volumes, you can use your FlashCopy repository more efficiently. Instead of requiring an equal amount of space to write data to, you can set aside a smaller amount of space in which to write data, where only the tracks that are changed are recorded. When your pool of storage is full, you can no longer perform a FlashCopy operation, and your session goes into a Severe state.

1. Follow these steps to create a FlashCopy session:
  - a. In the navigation tree, select **Sessions**.
  - b. Click **Create Session**.
  - c. In the Create Session wizard, for the **Choose Hardware Type** list, select the item that shows the type of storage system for the session.
  - d. In the **Choose Session Type** list, select **FlashCopy** and click **Next**.
  - e. On the Properties page, enter a session name, and a description. Complete the options that are described in "Session properties" on page 748. The options that are displayed depend on the storage system type. After you enter the information, click **Next**.
  - f. On the Site Locations page, select a location for Site 1 and click **Next**.
  - g. On the Results page, verify that the session was added successfully.
2. Follow these steps to add copy sets to the session:
  - a. On the Results page of the Create Session wizard, click **Launch Add Copy Sets Wizard**.
  - b. In the Add Copy Sets wizard, complete the following information. The field names that are displayed depend on the storage system type. When you have completed the information, click **Next**.

#### Storage system

Select a storage system. If the role has a location assigned to it, acceptable values for the storage system list are storage systems assigned to the same location as the role, and storage systems assigned to no location. In this case, the storage systems are grouped under different headings. If the role does not have a location, any storage system is acceptable.

**Logical storage system or I/O Group**

Select a logical storage system (LSS) or I/O group.

**Volume**

Select one volume or all volumes. The volumes are limited to the volumes within the LSS or I/O group that you selected.

**Session image**

Shows an image representing the session in which the role for which you are selecting volumes is highlighted. This image shows how many roles there are in the session and how the roles are distributed between the sites.

**Volume Details**

Shows information about the selected volume, including the volume name, full name, type, capacity, and whether the volume is protected and space efficient.

**Use a CSV file to import copy sets**

Select this option to import copy sets from a comma-separated value (CSV) file, click **Use a CSV file to import copy sets**. Type the full path name of the CSV file or click **Browse** to select the CSV file.

- c. On the Choose Target page, select the target storage system, LSS or I/O group, and volume. Click **Next**.
- d. On the Select Copy Sets page, select the copy sets that you want to add. You can click **Select All** to select all copy sets, **Deselect All** to select none of the copy sets, or **Add More** to add more copy sets to this session. Click **Next**.
- e. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.
- f. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.

**Creating a Snapshot session and adding copy sets**

A *snapshot* is a session type that creates a point-in-time copy of a volume or set of volumes without having to define a specific target volume. Snapshot sessions are available only for the XIV system.

1. Follow these steps to create a Snapshot session:
  - a. In the navigation tree, expand **Sessions**.
  - b. Click **Create Session**. The Create Session wizard starts.
  - c. In the **Choose Hardware Type** list, select **XIV**.
  - d. In the **Choose Session Type** list, select **Snapshot** and click **Next**.
  - e. On the Properties page, type a session name and description and click **Next**.
  - f. On the Site Locations page, select a location for Site 1 and click **Next**.
  - g. On the Results page, verify that the session was added successfully.
2. Follow these steps to add copy sets to the session:
  - a. On the Results page of the Create Session wizard, click **Launch Add Copy Sets Wizard**.
  - b. In the **Host1 storage system** list, select the storage system that contains the volumes that you want to add.

If the H1 role has an assigned location, only those storage systems that have the same location as the H1 role or storage systems that do not have a set location are displayed for selection.

If the H1 role does not have an assigned location, all storage systems are displayed for selection.

Storage systems that are assigned to a location are listed under the location name. Storage systems that are not assigned to a location are listed under the **None.** column.

- c. From the **Host1 pool** list, select the pool that contains the volumes.
- d. From the **Host1 volume** list, select the volumes. To select multiple volumes, press Ctrl or Shift and click the volumes in the list.
- e. If you want to import copy sets from a comma-separated value (CSV) file, click **Use a CSV file to import copy sets.** Type the full path name of the CSV file or click **Browse** to select the CSV file. Click **Next.**
- f. On the Matching Results page, click **Next** if the match was successful.
- g. On the Select Copy Sets page, select from the following options and click **Next.**

#### Select All

Click this button to select all of the copy sets in the table.

#### Deselect All

Click this button to clear all of the copy sets in the table.

#### Add More

Click this button to add another copy set to the list of copy sets to be created.

When you click **Add More**, you are returned to the Choose Host1 page of the wizard. The **Host1 storage system** and **Host1 pool** lists are populated with the values from the previously selected copy set. When you select the volumes for the various roles, the volumes are matched together and added to the list of copy sets on this page.

#### Selection check boxes

Select one or more copy sets that you want to create.

**Host 1** Lists the volume IDs that are associated with the Host1 role. You can click the link to display information about the volume, including the full name, type, capacity, and whether the volume is protected and space efficient.

#### Copy Set

Displays the copy set information for the specified copy sets and any warning or error messages that are associated with the copy set.

A warning or icon next to the **Show** button indicates that you cannot create a copy set for the H1 volume. Click **Show** to view the message.

- h. On the Confirm page, the number of copy sets to be added is displayed. Click **Next.**
- i. A progress bar is displayed. When the copy sets are added, review the results and click **Finish.**

## Creating a Metro Mirror session and adding copy sets

A Metro Mirror is a method of synchronous, remote data replication that operates between two sites that are up to 300 KM apart.

Follow the steps in this topic create a Metro Mirror session and add copy sets. If you want to enable HyperSwap or Open HyperSwap for the session, see “Managing a session with HyperSwap and Open HyperSwap replication” on page 100.

1. Follow these steps to create a Metro Mirror session:
  - a. In the navigation tree, select **Sessions**.
  - b. Click **Create Session**.
  - c. In the Create Session wizard, for the **Choose Hardware Type** list, select the item that shows the type of storage system for the session.
  - d. In the **Choose Session Type** list, select a Metro Mirror session type and click **Next**.
  - e. On the Properties page, enter a session name and description. Depending on the session type and storage system, additional options are displayed. Complete the options that are described in “Session properties” on page 748. After you enter the information on the Properties page, click **Next**.
  - f. On the site locations pages, choose a location for Sites 1 and 2 and click **Next**.
  - g. On the Results page, verify that the session was added successfully.
2. Follow these steps to add copy sets to the session:
  - a. On the Results page of the Create Session wizard, click **Launch Add Copy Sets Wizard**.
  - b. In the Add Copy Sets wizard, complete the following information for the Host pages. The field names that are displayed depend on the storage system type. When you have completed the information on each page, click **Next**.

#### **Storage system**

Select a storage system. If the role has a location assigned to it, acceptable values for the storage system list are storage systems assigned to the same location as the role, and storage systems assigned to no location. In this case, the storage systems are grouped under different headings. If the role does not have a location, any storage system is acceptable.

#### **Logical storage system, I/O Group, or Pool**

Select a logical storage system (LSS), I/O group, or pool.

#### **Volume**

Select one volume or all volumes. The volumes are limited to the volumes within the LSS, I/O group, or pool that you selected.

#### **Session image**

Shows an image representing the session in which the role for which you are selecting volumes is highlighted. This image shows how many roles there are in the session and how the roles are distributed between the sites.

#### **Volume Details**

Shows information about the selected volume, including the volume name, full name, type, capacity, and whether the volume is protected and space efficient.

#### **Use a CSV file to import copy sets**

Select this option to import copy sets from a comma-separated value (CSV) file, click **Use a CSV file to import copy sets**. Type the full path name of the CSV file or click **Browse** to select the CSV file.

- c. On the Select Copy Sets page, select the copy sets that you want to add. You can click **Select All** to select all copy sets, **Deselect All** to select none of the copy sets, or **Add More** to add more copy sets to this session. Click **Next**.
- d. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.
- e. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.

## Creating a Global Mirror session and adding copy sets

A Global Mirror is a method of asynchronous, remote data replication between two sites that are over 300 KM apart.

Follow the steps in this topic create a Global Mirror session and add copy sets.

1. Follow these steps to create a Global Mirror session:
  - a. In the navigation tree, select **Sessions**.
  - b. Click **Create Session**.
  - c. In the Create Session wizard, for the **Choose Hardware Type** list, select the item that shows the type of storage system for the session.
  - d. In the **Choose Session Type** list, select a Global Mirror session type and click **Next**.
  - e. On the Properties page, enter a session name and description. Depending on the session type and storage system, additional options are displayed. Complete the options that are described in “Session properties” on page 748. After you enter the information on the Properties page, click **Next**.
  - f. On the site locations pages, choose a location for Sites 1 and 2 and click **Next**.
  - g. On the Results page, verify that the session was added successfully.
2. Follow these steps to add copy sets to the session:
  - a. On the Results page of the Create Session wizard, click **Launch Add Copy Sets Wizard**.
  - b. In the Add Copy Sets wizard, complete the following information for the Host and Journal pages. The field names that are displayed depend on the storage system type. When you have completed the information on each page, click **Next**.

### Storage system

Select a storage system. If the role has a location assigned to it, acceptable values for the storage system list are storage systems assigned to the same location as the role, and storage systems assigned to no location. In this case, the storage systems are grouped under different headings. If the role does not have a location, any storage system is acceptable.

### Logical storage system, I/O Group, or Pool

Select a logical storage system (LSS), I/O group, or pool from this list.

### Volume

Select one volume or all volumes. The volumes are limited to the volumes within the LSS, I/O group, or pool that you selected.

You can use extent space-efficient volumes as copy set volumes for Global Mirror with Practice sessions for System Storage DS8000 6.3 or later. If you use an extent space-efficient volume as a source or

target volume in the copy set, all source and target volumes in the copy set must be extent space-efficient volumes. In this situation, the journal volumes can be extent space-efficient volumes, track space-efficient volumes, or a combination of both volume types. If you do not use an extent space-efficient volume as the source or target volume, journal volumes can be extent space-efficient, track space-efficient, and other types of volumes.

Extent space-efficient volumes must be fixed block (FB). You cannot use count key data (CKD) volumes.

#### **Session image**

Shows an image representing the session in which the role for which you are selecting volumes is highlighted. This image shows how many roles there are in the session and how the roles are distributed between the sites.

#### **Volume Details**

Shows information about the selected volume, including the volume name, full name, type, capacity, and whether the volume is protected and space efficient.

#### **Use a CSV file to import copy sets**

Select this option to import copy sets from a comma-separated value (CSV) file, click **Use a CSV file to import copy sets**. Type the full path name of the CSV file or click **Browse** to select the CSV file.

- c. On the Select Copy Sets page, select the copy sets that you want to add. You can click **Select All** to select all copy sets, **Deselect All** to select none of the copy sets, or **Add More** to add more copy sets to this session. Click **Next**.
- d. On the Confirm page, the number of copy sets to be added is displayed. Click **Next**.
- e. A progress bar is displayed. When the copy sets are added, review the results and click **Finish**.

### **Creating a Metro Global Mirror session and adding copy sets**

A Metro Global Mirror is a method of continuous, remote data replication that operates between three sites that varying distances apart. Metro Global Mirror combines Metro Mirror synchronous copy and Global Mirror asynchronous copy into a single session, where the Metro Mirror target is the Global Mirror source.

Follow the steps in this topic create a Metro Global Mirror session and add copy sets.

1. Follow these steps to create a Metro Global Mirror session:
  - a. In the navigation tree, select **Sessions**.
  - b. Click **Create Session**.
  - c. In the Create Session wizard, for the **Choose Hardware Type** list, select the item that shows the type of storage system for the session.
  - d. In the **Choose Session Type** list, select a Metro Global Mirror session type and click **Next**.
  - e. On the Properties page, enter a session name and a description. Complete the options that are described in "Session properties" on page 748. After you enter the information, click **Next**.
  - f. On the site locations pages, choose a location for Sites 1, 2, and 3 and click **Next**.





the Hardware Freeze Timeout Timer function within the storage system. If the controlling software loses connection to a storage system, the Metro Mirror relationships that it is controlling stay established and there is no way to *freeze* those pairs to create consistency across the multiple storage systems. When the freeze times out, dependent I/O is written to the target storage systems, which might corrupt data consistency. Freeze refers to a Metro Mirror (peer-to-peer remote copy [PPRC]) freeze function.

When determining whether to use the Metro Mirror heartbeat, analyze your business needs. Disabling the Metro Mirror heartbeat might result in data inconsistency. If you enable the Metro Mirror heartbeat and a freeze occurs, your applications will be unable to write during the freeze.

Metro Mirror heartbeat is disabled by default.

Metro Mirror heartbeat is not available for Metro Mirror with HyperSwap or Metro Global Mirror with HyperSwap.

There are two cases where lost communication between the coordination software (controller) and one or more storage systems can result in data consistency loss:

**Freeze event not detected by a disconnected storage system**

Consider a situation with four storage system machines in a primary site and four in a secondary site. One of the four storage systems on the primary loses the connection to the target site. This causes the affected storage system to prevent any writes from occurring, for a period determined by the Freeze timeout timer. At the same time, the affected storage controller loses communication with the controlling software and cannot communicate the Freeze event to the software.

Unaware of the problem, the controlling software does not issue the Freeze command to the remaining source storage systems. The freeze will stop dependent writes from being written to connected storage systems. However, once the Freeze times out and the long-busy is terminated, dependent write I/Os continue to be copied from the storage systems that did not receive the Freeze command. The Metro Mirror session is left in a state where one storage system has suspended copying while the other three storage systems are still copying data. This state causes inconsistent data on the target storage systems.

**Freeze event detected, but unable to propagate the Freeze command to all storage systems**

Consider a situation with four storage system machines in a primary site and four in a secondary site. One of the four storage systems on the primary loses the connection to the target site. This causes the affected storage system to issue long-busy to the applications for a period determined by the Freeze timeout timer. At the same time, one of the remaining three source systems loses communications with the controlling software.

The storage system that had an error writing to its target cannot communicate the Freeze event to the controlling software. The controlling software issues the Freeze command to all but the disconnected storage system (the one that lost communication with the software). The long-busy stops dependent writes from being written to the connected storage systems.

However, once the Freeze times out on the frozen storage system and the long-busy is terminated, dependent write I/Os continue to the target storage system from the source storage system that lost communication and did not receive the Freeze command. The Metro Mirror session is left a state where three storage systems have suspended copying and one storage system is still copying data. This state causes inconsistent data on the target storage systems.

Before IBM Tivoli Storage Productivity Center for Replication V3.1, if the controlling software within a Metro Mirror environment detected that a managed storage system lost its connection to its target, the controlling software stopped all the other source systems to ensure consistency across all the targets. However, if the controlling software lost communication with any of the source subsystems during the failure, it could not notify those storage systems of the freeze event or ensure data consistency. The Metro Mirror heartbeat helps to overcome this problem. In a high-availability configuration, the Metro Mirror heartbeat is continued by the standby server after the Takeover command is issued on the standby, enabling you to perform actions on the standby server without causing a freeze.

IBM Tivoli Storage Productivity Center for Replication registers with the managed ESS 800, DS6000 or DS8000 storage systems within a Metro Mirror session when the start command is issued to the session. After this registration occurs, a constant heartbeat is sent to the storage system. If the storage system does not receive a heartbeat from the IBM Tivoli Storage Productivity Center for Replication management server within the allotted time (a subset of lowest LSS timeout value across all the source LSSs), the storage system initiates a freeze. If IBM Tivoli Storage Productivity Center for Replication did not successfully communicate with the storage system, it initiates a freeze on the remaining storage system after the allotted time is expired.

**Note:** It is recommended that you avoid using the same LSS pairs for multiple Metro Mirror sessions. Metro Mirror uses a freeze command on ESS, DS6000, and DS8000 storage systems to create the data-consistent point. If there are other Metro Mirror sessions overlapping the same LSS pairs as in this session, those sessions are also suspended.

When you are using the Metro Mirror heartbeat, be aware that:

- The Metro Mirror heartbeat can cause a single point of failure: if an error occurs on just the management server and not the storage system, a freeze might occur.
- When the Metro Mirror heartbeat timeout occurs, the storage system remains in a long busy state for the duration of the LSS freeze timeout.

**Note:** If Metro Mirror heartbeat is enabled for storage systems that are connected through a HMC connection, a connection loss might cause lost Metro Mirror heartbeats, resulting in Freeze actions with application I/O impact for configured Extended Long Busy timeout.

The Metro Mirror heartbeat is supported on storage systems connected through a TCP/IP (direct connect or HMC) connection. It is not supported on storage systems connected through a z/OS connection. Enabling the Metro Mirror heartbeat with a z/OS connection does not fail; however, a warning message is displayed specifying that the Metro Mirror heartbeat function does not work unless you have an IP connection.

If Metro Mirror heartbeat is enabled for storage systems that are connected through a TCP/IP (either direct connect or HMC) connection and z/OS connection, and the TCP/IP connection fails, IBM Tivoli Storage Productivity Center for Replication suspends the Metro Mirror session because there is no heartbeat through the z/OS connection.

If Metro Mirror heartbeat is enabled for storage systems that are connected through a TCP/IP connection and z/OS connection and you remove all TCP/IP connections, IBM Tivoli Storage Productivity Center for Replication suspends the Metro Mirror sessions and the applications using those volume will be in Extended Long Busy timeout until the storage system's internal timeout timer expires. Ensure that you disable the Metro Mirror heartbeat for all Metro Mirror sessions before removing the last TCP/IP connection to avoid the Extended Long Busy timeout.

### **Enabling and disabling the Metro Mirror heartbeat**

The Metro Mirror heartbeat guarantees data consistency across multiple storage systems when the IBM Tivoli Storage Productivity Center for Replication management server cannot communicate with one or more storage systems. The Metro Mirror heartbeat is disabled by default.

To enable the Metro Mirror heartbeat, perform the following steps:

1. In the navigation tree, select **Advanced Tools**.
2. To enable the Metro Mirror heartbeat, click **Enable Heartbeat**.
3. To disable the Metro Mirror heartbeat, click **Disable Heartbeat**.

## **Exporting copy set data**

You can export data about all copy sets in a specific session, to maintain a backup copy that can be used to recover if you lose your session or upgrade to a different server.

Perform these steps to export the copy sets in a specific session:

1. In the navigation tree, select **Sessions**. The Sessions panel is displayed.
2. Select the session for which you want to export copy sets.
3. Select **Export Copy Sets** from the **Actions** list, and click **Go**. The Export Copy Set wizard displays the status of the export and a link to the exported file.
4. Click that link and save the file to the local system.

**Important:** You must save the file to your local system. After you close the panel, the data will be lost.

5. Click **Finish**.

## **Importing copy set data**

You can import copy set data that was previously exported to a comma separated value (CSV) file.

Perform the following steps to import copy sets into an existing session:

1. In the navigation tree, select **Sessions**. The Session panel is displayed.
2. Select the session for which you want to import copy sets.
3. Select **Add Copy Sets** from the **Actions** list, and click **Go**. The Add Copy Sets wizard is displayed.
4. Select **Use a CSV file to import copy sets**.

5. Type the location and name of the CSV file to import, or use **Browse** to select the file. Then, click **Next**.
6. Verify that the matching results were successful, and then click **Next**.
7. Select the copy sets you want to add, and then click **Next**.
8. Confirm the number of copy sets that you want to create, and click **Next**. A progress bar displays.
9. Click **Next**.
10. Verify the matches, and click **Finish**.

## Modifying the location of session sites

You can change the location associated with each site of in session.

**Prerequisites:** You must have Administrator privileges to modify the location of a site.

Changing the location of a site in a session does not affect the location of the storage systems associated with that site.

Perform these steps to modify the location of a site:

1. In the navigation tree, select **Sessions**. The Sessions panel is displayed.
2. Select the session whose site locations you want to change.
3. Click **Modify Site Location(s)** from the actions list, and then click **Go**. The Modify Site Locations wizard is displayed.
4. Select the location for the site from the drop-down list, and then click **Next**. Repeat for each available site.

To disable site awareness, set the location to None.

**Note:** You can select only locations that have already been associated with one or more storage systems.

5. Click **Next**.
6. Click **Finish**.

## Removing sessions

This topic describes how to remove sessions.

**Important:** You can remove only sessions that are in the Defined state.

Perform these steps to remove a session:

1. In the navigation tree, select **Sessions**.
2. Click the radio button next to the session you want to remove.
3. Select **Remove Session** from the **Actions** drop-down menu, and click **Go**.

## Removing copy sets

This topic describes how to remove copy sets.

Perform these steps to remove a copy set:

1. In the IBM Tivoli Storage Productivity Center for Replication navigation tree, select **Sessions**. Click the radio button next to the session that you want to remove copy sets from.

2. From the **Select Action** menu, select **Remove Copy Sets** and click **Go**. This starts the Remove Copy Sets wizard.
3. From the drop-down menus in the Remove Copy Sets wizard, select the Host 1 storage system, logical storage subsystem, and volume or select the all option. If you select all for a filter, the lower level filter or filters are disabled. Click **Next**.
4. Select the copy sets that you want to remove and click **Next**.
5. The number of copy sets to be removed is displayed. Select the following options for removing the copy sets and click **Next**:
  - **Do you want to keep the base relationships on the hardware, but remove the copy sets from the session?**
    - **Yes**. This option specifies that the base relationships remain on the hardware, but the copy sets are removed from the Tivoli Storage Productivity Center for Replication session. This option supports scenarios in which it might be best to leave the relationship on the hardware to avoid performing a full copy. For example, when you are migrating from one session type to another.  
Only the base relationships (Metro Mirror, Global Copy, Snapshot, and FlashCopy) are left on the hardware. The relationships are removed from any consistency groups that are defined on the storage system.
    - **No**. This option specifies that all relationships for the copy sets are removed from the hardware as well as the Tivoli Storage Productivity Center for Replication session. This option is the default.
  - **If there are errors removing relationships on the hardware, do you want to force the copy sets to be removed from the session?**
    - **Yes**. This option forces the removal of copy sets despite any errors that occur when removing the relationships from the storage system. Once a forced removal is complete, any relationships that remain on the storage system for that copy set must be removed manually using the storage system interface.
    - **No**. This option does not force the removal of copy sets. This option enables you to correct the errors and try to remove the copy sets again. This option is the default.
6. After the copy sets are removed, click **Finish**.

**Important:** If an application on the host has opened a device, the copy sets in a Tivoli Storage Productivity Center for Replication session for that device are removed, but the copy sets remain coupled on the host. To decouple the copy sets, see the Troubleshooting and support section of the Tivoli Storage Productivity Center for Replication Information Center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp>.

## Migrating an existing configuration to Tivoli Storage Productivity Center for Replication

This topic describes how to convert an existing DS6000 or DS8000 hardware relationship for a Metro Mirror, Global Mirror, or Metro Global Mirror session. This topic also describes how to assimilate pairs in a relationship for a Metro Mirror or Global Mirror session with SAN Volume Controller, Storwize V7000 Unified or Storwize V7000 volumes.

You can either complete this action manually or use the data migration utility for Tivoli Storage Productivity Center for Replication. You can download the data

migration utility from the Data Migration Utility for IBM(r) Tivoli(r) Storage Productivity Center for Replication. The data migration utility can handle all session types, and produces the Tivoli Storage Productivity Center for Replication CLI command script files and DSCLI script files necessary to migrate the relationships to Tivoli Storage Productivity Center for Replication.

Tivoli Storage Productivity Center for Replication cannot assimilate pairs that are suspended on the hardware. All pairs must be active for Tivoli Storage Productivity Center for Replication to assimilate the pairs.

### **Metro Mirror**

To manually assimilate pairs in a relationship for a Metro Mirror session, complete the following steps:

1. Determine which existing Metro Mirror pairs you want Tivoli Storage Productivity Center for Replication to manage, and record the names of your source and target volumes.
2. Create a Metro Mirror session in Tivoli Storage Productivity Center for Replication.
3. Using the Add Copy Set Wizard, choose the same source and target volumes that you identified in step 1.
4. Complete the Add Copy Set Wizard, and ensure that all copy sets are created successfully.
5. Issue a **Start** command from Tivoli Storage Productivity Center for Replication. At this time, the session automatically assimilates the relationships on the hardware (for example, if they are already in a Prepared state, Tivoli Storage Productivity Center for Replication will show them as Prepared).

**Note:** The **Start** command does not fully resynchronize data; instead, it detects the existing relationships and adopts them.

### **Global Mirror**

To manually assimilate pairs in a relationship for a Global Mirror session, complete the following steps:

1. Terminate the Global Mirror session on the hardware using the IBM System Storage DS command-line interface (DSCLI) or another application.
2. Remove the volumes from the Global Mirror session using the DSCLI or another application.
3. Close the session on each LSS upon which it is open using the DSCLI or another application.
4. Determine which existing Global Mirror pairs you want Tivoli Storage Productivity Center for Replication to manage, and record the names of your source and target volumes.
5. Create a Global Mirror session in Tivoli Storage Productivity Center for Replication.
6. Using the Add Copy Set Wizard, choose the same source and target volumes that you identified in step 4.
7. Complete the Add Copy Set Wizard, and ensure that all copy sets are created successfully.
8. Issue a **Start** command from Tivoli Storage Productivity Center for Replication. At this time, the session automatically assimilates the relationships on the hardware.

**Note:** The **Start** command does not fully resynchronize data; instead, it detects the existing relationships and adopts them.

### **Assimilating Metro Mirror pairs into a Three Site Metro Global Mirror session**

To manually assimilate a set of Metro Mirror pairs into a Three Site Metro Global Mirror session, complete the following steps:

1. Determine which existing Metro Mirror pairs you want Tivoli Storage Productivity Center for Replication to manage, and record the names of your source and target volumes.
2. Create a Metro Global Mirror session in Tivoli Storage Productivity Center for Replication.
3. Using the Add Copy Set Wizard, choose the same source and target volumes for the H1 and H2 roles of the Metro Global Mirror copy set, that you identified in step 1.
4. Complete the Add Copy Set Wizard, and ensure that all copy sets are created successfully.
5. Issue a **Start** command from Tivoli Storage Productivity Center for Replication. At this time, the session automatically assimilates the Metro Mirror relationships on the hardware (for example, if they are already in a Prepared state, Tivoli Storage Productivity Center for Replication will show them as Prepared). In addition, the **Start** command will establish the Global Copy relationships.

**Note:** The **Start** command does not fully resynchronize data on the Metro Mirror pairs assimilated; instead, it detects the existing relationships and adopts them. However, a full synchronization will be required for the Global Copy pairs.

### **Assimilating Three Site pairs into a Three Site Metro Global Mirror session**

To assimilate a full Three Site set of pairs into a Three Site Metro Global Mirror session, complete the following steps:

1. Terminate the Global Mirror session on the hardware using the IBM System Storage DS command-line interface (DSCLI) or another application.
2. Remove the volumes from the Global Mirror session using the DSCLI or another application
3. Close the session on each LSS upon which it is open using the DSCLI or another application.
4. Determine which existing Metro Mirror and Global Copy pairs you want Tivoli Storage Productivity Center for Replication to manage, and record the names of your source and target volumes.
5. Create a Metro Global Mirror session in Tivoli Storage Productivity Center for Replication.
6. Using the Add Copy Set Wizard, choose the same source and target volumes for the H1, H2 and H3 roles of the Metro Global Mirror copy set, that you identified in step 4.
7. Complete the Add Copy Set Wizard, and ensure that all copy sets are created successfully.
8. Issue a **Start** command from IBM Tivoli Storage Productivity Center for Replication. At this time, the session automatically assimilates the Metro Mirror and Global Copy relationship on the hardware.



**Note:** The **Start** command does not fully resynchronize data on the Metro Mirror and Global Copy pairs assimilated; instead, it detects the existing relationships and adopts them.

### **Global Mirror and Metro Mirror assimilation for SAN Volume Controller, Storwize V7000, Storwize V7000 Unified or the XIV system**

To manually assimilate pairs in a relationship for a Metro Mirror or Global Mirror session with SAN Volume Controller, Storwize V7000, Storwize V7000 Unified or the XIV system volumes, complete the following steps:

1. Remove all pairs from the consistency group and delete the consistency group on the storage system using CLI for the storage system.
2. Determine which existing Metro Mirror and Global Mirror pairs you want Tivoli Storage Productivity Center for Replication to manage, and record the names of your source and target volumes.
3. Create a Metro Mirror or Global Mirror session in Tivoli Storage Productivity Center for Replication.
4. Using the Add Copy Set Wizard, choose the same source and target volumes that you identified in step 2.
5. Complete the Add Copy Set Wizard, and ensure that all copy sets are created successfully.
6. Issue a **Start** command from Tivoli Storage Productivity Center for Replication. The session automatically assimilates the relationships on the hardware and adds them to a new consistency group on the storage system.

**Note:** The **Start** command does not fully resynchronize data; instead, it detects the existing relationships and adopts them placing them in a consistency group managed by the session.

---

## **Practicing disaster recovery**

You can use practice volumes to test your disaster recovery actions while maintaining disaster recovery capability. Practice volumes are available in Metro Mirror Failover and Failback sessions, Global Mirror Failover and Failback, Global Mirror Either Direction sessions, and Metro Global Mirror with Practice sessions.

**Note:** You can test your disaster recovery actions without using practice volumes; however, without practice volumes, you cannot maintain disaster recovery capability while continuing to copy the data.

### **Practice volumes**

You can use a *practice volume* to practice what you would do in the event of a disaster, without interrupting current data replication. Practice volumes are available in Metro Mirror, Global Mirror, and Metro Global Mirror sessions.

To use the practice volumes, the session must be in the prepared state. Issuing the Flash command against the session while in the Prepared state creates a usable practice copy of the data on the target site.

**Note:** You can test disaster-recovery actions without using practice volumes; however, without practice volumes, you cannot continue to copy data changes between volumes while testing disaster-recovery actions.

## Practicing disaster recovery for a Metro Mirror Failover/Failback with Practice session

A Metro Mirror Failover and Failback session with Practice combines Metro Mirror and FlashCopy to provide a point-in-time copy of the data on the remote site. You can use this to practice what you might do if a disaster occurred, without losing your disaster recovery capability.

This function is available on the following storage systems:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified
- TotalStorage Enterprise Storage Server Model 800
- System Storage DS8000
- System Storage DS6000

Perform these steps to practice disaster recover actions for a Metro Mirror Failover/Failback with Practice session:

1. Start a Metro Mirror with Practice session.
2. When the Metro Mirror session reaches the Prepared state, issue a Flash command to make a point-in-time copy of the data on H2. This creates a consistent point-in-time copy of your data on the H2 volume and then restarts the session so the copying from H1 to I2 continues. This temporarily stops copying of the data from site 1 to site 2, and creates a consistent point-in-time copy. The data replication from H1 to I2 is then restarted.

**Note:** For ESS, DS6000, DS8000 storage systems, the **Flash** command uses the freeze and thaw processing to create a data consistent point for the FlashCopy. If there is another Metro Mirror session overlapping on one or more of the same LSS pairs, that session will be suspended. It is also possible that the suspension of the other session might cause the Metro Mirror session to remain suspended after the flash command is issued instead of returning to Prepared state. Avoid using the same LSS pairs for multiple Metro Mirror sessions if possible.

3. Practice the same actions you would take in an actual disaster, using H2 as your new host volume, while running the real application on H1.

## Practicing disaster recovery for a Global Mirror either Direction with two-site Practice session

A Global Mirror (either direction) with two-site Practice combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on a remote site at a distance over 300 km away from your first site. You can use this to practice what you might do if a disaster occurred, without losing your disaster recovery capability.

**Note:** This function is available only on ESS, DS6000, and DS8000 storage systems.

Perform these steps to practice disaster recover actions for a Global Mirror either Direction with two-site Practice session:

1. Start a Global Mirror with Practice session.
2. When the session reaches the Prepared state, issue a Flash command to restore consistent data on I2 and make a point-in-time copy of the data on H2. This creates a consistent point-in-time copy of your data on the H2 volume and then

restarts the session so the copying from H1 to I2 continues. This temporarily stops copying of the data from site 1 to site 2, and creates a consistent point-in-time copy. The data replication from H1 to I2 is then restarted.

**Note:** FlashCopy must be always a full copy due to limitations of the hardware.

3. Practice the same actions you would take in an actual disaster, using H2 as your new host volume, while running the real application on H1.

**Note:** With two directions, you can reverse the direction of your data flow.

## Practicing disaster recovery for a Global Mirror Failover/Failback with Practice session

A Global Mirror Failover and Failback with Practice combines Global Mirror and FlashCopy to provide a point-in-time copy of the data on a remote site at a distance over 300 km away from your first site. You can use this to practice what you might do if a disaster occurred.

You can do this practice without losing your disaster recovery capability. The number of volumes used for the device varies, but the steps to conduct a Global Mirror Failover and Failback with Practice are the same for both devices.

This function is available on the following storage systems:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified
- TotalStorage Enterprise Storage Server Model 800
- System Storage DS8000
- System Storage DS6000

Perform these steps to practice disaster recover actions for a Global Mirror Failover/Failback with Practice session:

1. Start a Global Mirror with Practice session.
2. When the session reaches the Prepared state, issue a Flash command to restore consistent data on I2 and make a point-in-time copy of the data on H2. This creates a consistent point-in-time copy of your data on the H2 volume and then restarts the session so the copying from H1 to I2 continues. This temporarily stops copying of the data from site 1 to site 2, and creates a consistent point-in-time copy. The data replication from H1 to I2 is then restarted.

**Note:** FlashCopy must be always a full copy due to limitations of the hardware.

3. Practice the same actions you would take in an actual disaster, using H2 as your new host volume, while running the real application on H1.

## Practicing disaster recovery for a Metro Global Mirror Failover/Failback with Practice session

A Metro Global Mirror Failover/Failback with Practice session combines Metro Mirror, Global Mirror and FlashCopy across three sites to provide a point-in-time copy of the data on the third site. You can use this to practice what you might do if a disaster occurred without losing your disaster recovery capability.

**Note:** This function is available on ESS, DS6000, and DS8000 storage systems.

The intermediate volume is on the third site (I3). This maintains disaster recovery capability while a copy is kept on the H3 volume for practice purposes.

Perform these steps to practice disaster recover actions for a Metro Global Mirror Failover/Failback with Practice session:

1. Start a Metro Global Mirror with Practice session.
2. When the session reaches the Prepared state, issue a Flash command to take a point-in-time copy of the data that is on I3, on H3. This creates a consistent point-in-time copy of your data on the H3 volume, and then automatically restarts the session so that copying from H1 to H2 to I3 continues. The Flash command temporarily stops copying the data from site 2 to site 3, in order to create a consistent point-in-time copy on I3, while maintaining disaster recovery capabilities on site 2 using the Metro Mirror portion of the session. Then, data replication from H2 to I3 is restarted.

**Note:** FlashCopy must be always a full copy due to hardware limitations.

3. Practice the same actions you would take in an actual disaster, using H3 as your practice host volume, while you run the real application on H1. This enables you to use the same scripts and commands to run on H3 that you would use in an actual disaster.

---

## Monitoring health and status

### Viewing the health summary

Use the Health Overview panel to view overall health and status of sessions, storage systems, host systems, and management servers.

The Health Overview panel is the first panel that you see after you log on. You can display this panel by selecting **Health Overview** in the IBM Tivoli Storage Productivity Center for Replication navigation tree. This panel provides the following information:

- Overall session status: indicates session status, which can be normal, warning, or severe. The status can also be inactive, if all sessions are defined or if no sessions exit.
- Overall storage system status: indicates the connection status of storage systems.
- Overall host system status: indicates the connection status of host systems.
- Management server status (applicable only if you are using a Business Continuity license): indicates the status of the standby server if you are logged on to the local server. If you are logged on to the standby server, this status indicates the status of the local server.

Health information is always shown as a mini-panel under the navigation tree.

### Viewing SNMP alerts

IBM Tivoli Storage Productivity Center for Replication SNMP trap descriptions can be viewed from the IBM Tivoli Storage Productivity Center Alert panel.

From the IBM Tivoli Storage Productivity Center GUI, in the navigation tree, expand **Alerting > Alert Log > All** and select **Replication**.

## Viewing sessions





This section describes how to view sessions details and session properties.

### Session status icons

The IBM Tivoli Storage Productivity Center for Replication GUI uses icons to represent the status of each session.

The following table describes each session status icon.

Table 47. Session status icons

| Icon                                                                               | Meaning  | Description                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------------------------------------------------------------|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|   | Inactive | The session is in a defined state, with no activity on the hardware.                                                                                                                                                                                                                                                                                                                                                   |
|   | Normal   | A consistent copy of the data either exists or is being maintained.                                                                                                                                                                                                                                                                                                                                                    |
|   | Warning  | For Metro Mirror, Global Mirror, and Metro Global Mirror, the session might have volumes that are being synchronized or are about to be synchronized, with no suspended volumes. For FlashCopy, the Warning status is valid only after the start command is issued and before the flash, and means that the session is either preparing or is ready for a flash command but targets do not yet have a consistent copy. |
|  | Severe   | One or more errors must be dealt with immediately. Possible causes include the following: <ul style="list-style-type: none"> <li>• One or more volumes are suspended</li> <li>• A session is suspended</li> <li>• A volume is not copying correctly</li> </ul>                                                                                                                                                         |

### Session images

The IBM Tivoli Storage Productivity Center for Replication GUI provides a visual aid to help you create and manage your sessions. The visual aid shows the number of volume roles in the session and how the roles are distributed between the sites. It also shows the copy method and direction.

### Volume role symbols

The volume role symbols represent the replication status on the volumes.

Table 48. Volume role symbols







| Symbol                                                                              | Description          | Meaning                                                                                                                                                  |
|-------------------------------------------------------------------------------------|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | Active host volumes  | This symbol represents volumes that contain the source of updated tracks to which the application is actively issuing read and write input/output (I/O). |
|  | Recoverable volumes  | This symbol represents volumes that contain a consistent copy of the data.                                                                               |
|  | Inconsistent volumes | This symbol represents the volumes that do not contain a consistent copy of the data.                                                                    |

Table 48. Volume role symbols (continued)

| Symbol                                                                            | Description      | Meaning                                                                                                                                     |
|-----------------------------------------------------------------------------------|------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
|  | Selected volumes | This symbol represents the volumes that are selected for an operation (for example, changing location or displaying role pair information). |
|  |                  |                                                                                                                                             |
|  |                  |                                                                                                                                             |

## Data copying symbols

The data copying symbols are arrows indicate the type of copy that occurs between the volume roles. The direction of the arrow indicates the direction of the copy.

Table 49. Data copying symbols


















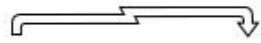



| Symbol                                                                              | Description                                                                                  | Meaning                                                                                                                                               |
|-------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
|    | FlashCopy copying (the lightning bolt indicates the direction of the FlashCopy)              | This symbol represents a FlashCopy relationship in which data is being copied from the host to the target.                                            |
|    |                                                                                              |                                                                                                                                                       |
|  | FlashCopy copying with errors (the lightning bolt indicates the direction of the FlashCopy)  | This symbol represents a FlashCopy relationship in which data is being copied from the host to the target, but there are errors on one or more pairs. |
|  |                                                                                              |                                                                                                                                                       |
|  | FlashCopy inactive (the lightning bolt indicates the direction of the FlashCopy)             | This symbol represents an inactive FlashCopy relationship.                                                                                            |
|  |                                                                                              |                                                                                                                                                       |
|  | FlashCopy inactive with errors (the lightning bolt indicates the direction of the FlashCopy) | This symbol represents an inactive FlashCopy relationship in which there are errors on one or more pairs.                                             |
|  |                                                                                              |                                                                                                                                                       |
|  | Metro Mirror copying                                                                         | This symbol represents a copying Metro Mirror relationship.                                                                                           |
|  | Metro Mirror copying with errors                                                             | This symbol represents a Metro Mirror relationship that is copying, but with errors on one or more pairs.                                             |
|  | Metro Mirror inactive                                                                        | This symbol represents an inactive Metro Mirror relationship.                                                                                         |

Table 49. Data copying symbols (continued)

| Symbol                                                                              | Description                       | Meaning                                                                                                                                                                                                                                                                                       |
|-------------------------------------------------------------------------------------|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|    | Metro Mirror inactive with errors | This symbol represents an inactive Metro Mirror relationship with errors on one or more pairs.                                                                                                                                                                                                |
|    | Global Copy copying               | This symbol represents a copying Global Copy relationship.                                                                                                                                                                                                                                    |
|    |                                   |                                                                                                                                                                                                                                                                                               |
|    | Global Copy copying with errors   | This symbol represents a copying Global Copy relationship with errors on one or more pairs.                                                                                                                                                                                                   |
|    |                                   |                                                                                                                                                                                                                                                                                               |
|    | Global Copy inactive              | This symbol represents an inactive Global Copy relationship.                                                                                                                                                                                                                                  |
|    |                                   |                                                                                                                                                                                                                                                                                               |
|    | Global Copy inactive with errors  | This symbol represents an inactive Global Copy relationship with errors on one or more pairs.                                                                                                                                                                                                 |
|  |                                   |                                                                                                                                                                                                                                                                                               |
|  | HyperSwap or Open HyperSwap       | This symbol HyperSwap or Open HyperSwap for a session. If a failure occurs when input/output (I/O) is being written to the primary storage system, HyperSwap or Open HyperSwap automatically swap the I/O to the secondary site with no user interaction and little or no application impact. |

## Session states

You can view the health and status of a session in the Tivoli Storage Productivity Center for Replication GUI.

**Important:** Use only the Tivoli Storage Productivity Center for Replication GUI or CLI to manage session relationships, such as volume pairs and copy sets. Do not modify session relationships through individual hardware interfaces, such as the DSCLI. Modifying relationships through the individual hardware interfaces can result in a loss of consistency across the relationships managed by the session, and might cause the session to be unaware of the state or consistency of the relationships.

You can use the Refresh States command is used to refresh the states of the session. Issue this command to query the states of the copy sets on the hardware. You do not need to run this command under typical circumstances; Tivoli Storage Productivity Center for Replication refreshes the states of its sessions through

multiple means. However, if you discover an inconsistency between Tivoli Storage Productivity Center for Replication and the hardware, you can use this command to enable IBM Tivoli Storage Productivity Center for Replication to update itself. Because this command triggers multiple queries on the hardware (having an adverse impact on hardware performance if you run it too often), you can only execute this command every few minutes in each session.

The following table describes each session state.

*Table 50. Session states*

| State    | Session type | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Defined  | All          | The session exists but is inactive.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Flashing | All          | <p>In a Metro Mirror or Global Mirror session, data copying is temporarily suspended while a consistent practice copy of data is being prepared on site 2.</p> <p>In a Metro Global Mirror session, data copying is temporarily suspended while a consistent practice copy of data is being prepared on site 3.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Prepared | All          | <p>The source to target data transfer is active.</p> <p>In a Metro Mirror, Global Mirror, or Metro Global Mirror session, the data written to the source is transferred to the target, and all volumes are consistent and recoverable.</p> <p>In a FlashCopy session, the volumes are not yet consistent, but the flash is ready to begin.</p> <p><b>Note:</b> For sessions on the following storage systems, do not alter the relationships on the hardware that you established with Tivoli Storage Productivity Center for Replication:</p> <ul style="list-style-type: none"> <li>• SAN Volume Controller</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> </ul> <p>For example, if a Metro Mirror session with one copy set is in the Prepared state, and you stop the role pair, the session still displays in the Prepared state.</p> |



Table 50. Session states (continued)

| State             | Session type        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Preparing         | All                 | <p>The volumes are initializing, synchronizing, or resynchronizing.</p> <p>In a Metro Mirror, Global Mirror, or Metro Global Mirror session, synchronization occurs after the first <b>Start</b> command is issue on a session. Resynchronization occurs when a volume was prepared and then suspended. The hardware records the changed tracks so that on the next startup, only the changed tracks are copied.</p> <p>In a FlashCopy session, the volumes are initializing. The preparing state for FlashCopy sessions applies only to the following storage systems:</p> <ul style="list-style-type: none"> <li>• SAN Volume Controller</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> </ul> |
| Recovering        | All                 | The session is in the process of recovering.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Suspended         | All                 | <p>Data copying has temporarily stopped.</p> <p><b>Note:</b> The suspended state applies only to Global Mirror, Metro Mirror, and Metro Global Mirror sessions.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| SuspendedH1H2 MGM | Metro Global Mirror | Data copying between site 1 and site 2 is suspended.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| SuspendedH1H3 MGM | Metro Global Mirror | Data copying between site 1 and site 3 is suspended.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Suspending        | All                 | <p>The session is transitioning into a Suspended state.</p> <p><b>Note:</b> The Suspending state applies only to Global Mirror and Metro Global Mirror sessions and does not apply to the following storage systems:</p> <ul style="list-style-type: none"> <li>• SAN Volume Controller</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> </ul>                                                                                                                                                                                                                                                                                                                                                    |
| Target available  | All                 | Target volumes are available for application updates.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

Table 50. Session states (continued)

| State       | Session type | Description                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Terminating | FlashCopy    | <p>The session is being terminated because you issued a Terminate action under the following conditions:</p> <ul style="list-style-type: none"> <li>You permitted the target to be Metro Mirror or Global Copy.</li> <li>You set the Require or Attempt to Preserve Mirror option.</li> </ul> <p>The session displays as Terminating until the FlashCopy background copy is complete and no longer exists on the hardware.</p> |

## Session properties

This topic lists the session properties.

Session properties include the following attributes:

### Description

Type the description for this session.

### ESS/DS FlashCopy Options

Select the FlashCopy options that you want to associate with this session:

#### Incremental

Select this option to set up the relationship for recording changes to the H1 volume. Any subsequent FlashCopy operation for that session copies only the tracks that have changed since the last flash. Incremental always assumes persistence.

**Note:** This option is used along with the Persistent and No-Copy options.

#### Persistent

Select this option to keep the relationship established on the hardware after all source tracks are copied to the target. If this option is not selected, the local replication relationship ends after the T1 volume contains a complete point-in-time image of the H1 volume.

**Note:** If you select the Incremental option, the Persistent option is automatically selected as well.

#### No Copy

Select this option if you do not want the hardware to write the background copy until the source track is written to. Replication is done by using a copy-on-write technique. This technique does not copy data to the T1 volume until the blocks or tracks of the H1 volume are modified. The point-in-time volume image is composed of the unmodified data on the H1 volume, and the data that was copied to the T1 volume. If you want a complete point-in-time

copy of the H1 volume to be created on the T1 volume, do not use the No Copy option. This option causes the data to be asynchronously copied from the H1 volume to the T1 volume.

**Note:** Although you can select any space-efficient volume as the target, you cannot change the Permit Space Efficient Target flag. This flag is always set. When selecting space-efficient volumes as targets, you might receive an x0FBD error message if you attempt a full background copy. To avoid this message, select the No Copy option.

**Allow FlashCopy target to be Metro Mirror source**

Select this option if the target of the FlashCopy relationship is to overwrite all data on the source of a Metro Mirror relationship. If this option is cleared, a flash to a Metro Mirror source volume fails.

If you select this option, you must also select one of the following policies:

**Do not attempt to preserve Metro Mirror consistency**

Select this option if the Metro Mirror or Global Copy pair at the target of the FlashCopy relationship is to perform a full copy of the data to the secondary of the Metro Mirror or Global Copy pair.

**Attempt to preserve Metro Mirror consistency, but allow FlashCopy even if Metro Mirror target consistency cannot be preserved**

Select this option to attempt to preserve the consistency of the Metro Mirror relationship at the target of the FlashCopy relationship when both the source and target of the FlashCopy relationship are a source of a Metro Mirror relationship. If the consistency cannot be preserved, a full copy of the Metro Mirror relationship at the target of the FlashCopy relationship is performed. To preserve consistency, parallel flashes are performed (if possible) on both sites.

**Note:** This option is available only on IBM System Storage DS8000 storage systems with the required code levels installed.

**Attempt to preserve Metro Mirror consistency, but fail FlashCopy if Metro Mirror target consistency cannot be preserved**

This option prevents a full copy from being performed over the Metro Mirror link. Instead, parallel flashes are performed (if possible) on both sites. If the consistency cannot be preserved, the Flash for the FlashCopy relationships fails, and the data of the Metro Mirror relationship at the target of the FlashCopy relationship is not changed.

**Note:** This option is available only on System Storage DS8000 storage systems with the required code levels installed.

## **SAN Volume Controller / Storwize V7000 / Storwize V7000 Unified FlashCopy Options**

Select the FlashCopy options that you want to associate with this session.

### **Incremental**

Select this option to set up the relationship for recording changes to the H1 volume. Any subsequent FlashCopy operation for that session copies only the tracks that have changed since the last flash. Incremental always assumes persistence.

### **Background Copy Rate**

Type the copy rate that the Storage System uses to perform the background copy of the FlashCopy role pair. You can specify a percentage between 0 and 100. The default is 50. Specifying 0 is equivalent to specifying the No-Copy option on a TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, or System Storage DS6000 FlashCopy session.

You can modify this value at any time during the session. If the session is performing a background copy when you change the option, Tivoli Storage Productivity Center for Replication immediately modifies the background copy rate of the consistency group on the storage system. The storage system consistency group immediately starts by using this new rate to complete the background copy.

## **Basic HyperSwap Options**

Select the Basic HyperSwap options that you want to associate with this session.

If Tivoli Storage Productivity Center for Replication is running on z/OS, and volumes are attached by a fibre-channel connection, Tivoli Storage Productivity Center for Replication can manage a HyperSwap session. With this option selected, a failure on the host accessible volumes triggers a HyperSwap, redirecting application input/output (I/O) to the secondary volumes.

These options are available only for Basic HyperSwap sessions:

### **Disable HyperSwap**

Select this option to prevent a HyperSwap from occurring by command or event.

### **On Configuration Error:**

#### **Partition the system(s) out of the sysplex**

Select this option to partition out of the sysplex when a new system is added to the sysplex and encounters an error in loading the configuration. Restart the system if you select this option.

#### **Disable HyperSwap**

Select this option to prevent a HyperSwap from occurring by command or event.

### **On Planned HyperSwap Error:**

#### **Partition out the failing system(s) and continue swap processing on the remaining system(s)**

Select this option to partition out the failing system and continue the swap processing on any remaining systems.

### **Disable HyperSwap after attempting backout**

Select this option to stop the HyperSwap action, and disable the HyperSwap commands or events.

### **On Unplanned HyperSwap Error:**

#### **Partition out the failing system(s) and continue swap processing on the remaining system(s)**

Select this option to partition out the failing systems and continue HyperSwap processing on the remaining systems when a new system is added to the sysplex and HyperSwap does not complete. You must restart the system if you select this option.

#### **Disable HyperSwap after attempting backout**

Select this option to stop the HyperSwap action, and disable the HyperSwap commands or events.

### **Fail MM/GC if target is online (CKD only)**

Select this option to fail any session commands for a copy set if the session is online and if the target volume in the copy set is visible to a host.

### **ESS/DS Metro Mirror Options**

Select the Metro Mirror options that you want to associate with this session. These options are available only for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, or System Storage DS6000 storage systems in a Metro Mirror session:

#### **Reset Secondary Reserve**

Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets being started when a Start command is issued for the session.

#### **Note:**

Before enabling the Reset Secondary Reserves option, be aware that this action causes the session to overwrite all data on the target volume.

### **Fail MM/GC if target is online (CKD only)**

Select this option to fail any session commands for a copy set if the session is online and if the target volume in the copy set is visible to a host.

### **Metro Mirror Suspend Policy**

You must select either **Hold I/O after Suspend** or **Release I/O after Suspend** when a replication link suspends. A replication link suspends for various reasons. In some cases, it happens because the replication link is down or the auxiliary storage controller has failed. In other cases, it might be the first failure in an eventual full failure at the primary site.

If your primary goal is zero data loss, specify **Hold I/O after Suspend**. Selecting this option means that no additional updates are made to storage until the entire scope of the failure can be determined. If you select this option, the systems by using that storage become unavailable until a decision is made about how to proceed. It might also require restarting the systems.

If your primary goal is availability, specify **Release I/O after Suspend**. This option enables workloads to continue running after all the replication links have been suspended, maintaining a consistent copy of data at the remote site. If the failure spreads through the primary site after the workloads resume running and results in a site failure, data loss occurs. The I/O requests that are issued after the resume and before the failure at the primary site are lost. The storage at the secondary site remains consistent, however, and you can recover it there.

#### **Hold I/O after Suspend**

Select this option to block the application from writing while a consistent copy of the data is maintained on the remote site; however, it does not automatically release the application. This option keeps the source equal to the target. You must use the **Release I/O** command on the session or wait for the Hardware Freeze Timeout timer to expire before the application can continue to write to the source.

#### **Release I/O after Suspend**

Select this option to block writes to the application while a consistent copy of the data is formed on the remote site. This operation is followed immediately by releasing the block so that the application can continue writing to the source. This option allows for little application impact, but causes the source to potentially be different from the target.

This option is the default setting for all new sessions.

#### **DS FlashCopy Options for Role pair H2-I2**

The following option is available only for System Storage DS8000 version 4.2 or later.

##### **Persistent**

Select this option to keep FlashCopy pairs persistent on the hardware.

#### **Manage H1-H2 with HyperSwap**

Select this option to trigger a HyperSwap, redirecting application I/O to the secondary volumes, when there is a failure on the host accessible volumes. For Metro Global Mirror sessions, the Global Mirror portion of the session continues to run uninterrupted.

Tivoli Storage Productivity Center for Replication can manage the H1-H2 sequence of a Metro Mirror or Metro Global Mirror session by using HyperSwap if Tivoli Storage Productivity Center for Replication is running on z/OS, and volumes are attached by a fibre-channel connection.

##### **Notes:**

- When this option is selected, the **Suspend H1-H2** command is available only if the Disable HyperSwap option is also selected.
- Setting this option automatically sets the Release IO after suspend Metro Mirror policy.

##### **Disable HyperSwap**

Select this option to prevent a HyperSwap from occurring by command or event.

### **On Configuration Error:**

#### **Partition the system(s) out of the sysplex**

Select this option to partition out of the sysplex when a new system is added to the sysplex and encounters an error in loading the configuration. Restart the system if you select this option.

#### **Disable HyperSwap**

Select this option to prevent a HyperSwap from occurring by command or event.

### **On Planned HyperSwap Error:**

#### **Partition out the failing system(s) and continue swap processing on the remaining system(s)**

Select this option to partition out the failing system and continue the swap processing on any remaining systems.

#### **Disable HyperSwap after attempting backout**

Select this option to stop the HyperSwap action, and disable the HyperSwap commands or events.

### **On Unplanned HyperSwap Error:**

#### **Partition out the failing system(s) and continue swap processing on the remaining system(s)**

Select this option to partition out the failing systems and continue HyperSwap processing on the remaining systems when a new system is added to the sysplex and HyperSwap does not complete. You must restart the system if you select this option.

#### **Disable HyperSwap after attempting backout**

Select this option to stop the HyperSwap action, and disable the HyperSwap commands or events.

### **Manage H1-H2 with Open HyperSwap**

If volumes are attached to an IBM AIX host, Tivoli Storage Productivity Center for Replication can manage the H1-H2 sequence of a Metro Mirror session by using Open HyperSwap. If this option is selected, a failure on the host accessible volumes triggers a swap, which redirects application I/O to the secondary volumes. Only volumes that are currently attached to the host systems that are defined on the Tivoli Storage Productivity Center for Replication Host Systems panel are eligible for Open HyperSwap.

#### **Disable Open HyperSwap**

Select this option to prevent a swap from occurring by a command or event while keeping the configuration on the host system and all primary and secondary volumes coupled.

### **SAN Volume Controller / Storwize V7000 / Storwize V7000 Unified Metro Mirror Options**

Select the Metro Mirror options that you want to associate with this session. The following options are available only for storage systems that are in a Metro Mirror Failover/Failback with Practice session:

**Incremental**

Select this option to set up the relationship for recording changes to the practice volume (H2). All subsequent FlashCopy operations between the intermediate volume and the host volume copy only the data that has changed since the previous FlashCopy operation. Incremental always assumes persistence.

**Background Copy Rate for H2-I2**

Type the copy rate that the Storage System uses to perform the background copy of the FlashCopy role pair. You can specify a percentage between 0 and 100. The default is 50. Specifying 0 is equivalent to specifying the No-Copy option on a TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, or System Storage DS6000 FlashCopy session.

You can modify this value at any time during the session. If the session is performing a background copy when you change the option, Tivoli Storage Productivity Center for Replication immediately modifies the background copy rate of the consistency group on the storage system. The storage system consistency group immediately starts by using this new rate to complete the background copy.

**ESS/DS Global Mirror Options**

Select the Global Mirror options that you want to associate with this session. These options are available only for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, and System Storage DS6000 storage systems in a Global Mirror session:

**Consistency group interval time (seconds)**

Type how often, in seconds, the Global Mirror session attempts to form a consistency group. When you reduce this value, it might be possible to reduce the data exposure of the session. A lower value causes the session to attempt to create consistency groups more frequently, which might affect the storage systems (for example, by increasing the processing load and message traffic load).

**Reset Secondary Reserve**

Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets being started when a Start command is issued for the session.

**Note:**

Before enabling the Reset Secondary Reserves option, be aware that this action causes the session to overwrite all data on the target volume.

**Fail MM/GC if target is online (CKD only)**

Select this option to fail any session commands for a copy set if the session is online and if the target volume in the copy set is visible to a host.

**DS FlashCopy Options for Role pair H2-I2**

The following option is available only for System Storage DS8000 version 4.2 or later.

**Persistent**

Select this option to keep FlashCopy pairs persistent on the hardware.



### **No Copy**

Select this option if you do not want the hardware to write the background copy until the source track is written to. Replication is done by using a copy-on-write technique. This technique does not copy data to the I2 volume until the blocks or tracks of the H2 volume are modified. The point-in-time volume image is composed of the unmodified data on the H2 volume, and the data that was copied to the I2 volume. If you want a complete point-in-time copy of the H2 volume to be created on the I2 volume, do not use the No Copy option. This option causes the data to be asynchronously copied from the H2 volume to the I2 volume.

**Note:** Although you can select any space-efficient volume as the target, you cannot change the Permit Space Efficient Target flag. This flag is always set. When selecting space-efficient volumes as targets, you might receive an x0FBD error message if you attempt a full background copy. To avoid this message, select the No Copy option.

### **DS FlashCopy Options for Role pair I2-J2**

The following option is available only for System Storage DS8000 version 4.2 or later.

#### **Reflash After Recover**

Select this option if you want to create a FlashCopy replication between the I2 and J2 volumes after the recovery of a Global Mirror session. If you do not select this option, a FlashCopy replication is created only between the I2 and H2 volumes.

### **SAN Volume Controller / Storwize V7000 / Storwize V7000 Unified Global Mirror Options**

Select the Global Mirror options that you want to associate with this session. These options are available only for storage systems that are in a Global Mirror Failover/Failback with Practice session:

#### **Incremental**

Select this option to set up the relationship for recording changes to the practice volume (H2). All subsequent FlashCopy operations between the intermediate volume and the host volume copy only the data that has changed since the previous FlashCopy operation. Incremental always assumes persistence.

#### **Background Copy Rate for H2-I2**

Type the copy rate that the Storage System uses to perform the background copy of the FlashCopy role pair. You can specify a percentage between 0 and 100. The default is 50. Specifying 0 is equivalent to specifying the No-Copy option on a TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, or System Storage DS6000 FlashCopy session.

You can modify this value at any time during the session. If the session is performing a background copy when you change the option, Tivoli Storage Productivity Center for Replication immediately modifies the background copy rate of the consistency

group on the storage system. The storage system consistency group immediately starts by using this new rate to complete the background copy.

#### **XIV Global Mirror Options**

Select the Global Mirror options that you want to associate with this session. These options are available only for the XIV systems in a Global Mirror Failover/Failback session:

##### **Recovery point objective threshold (seconds)**

Type the number of seconds that you want to set for the recovery point objective (RPO) threshold. RPO represents a measure of the maximum data loss that is acceptable in the event of a failure or unavailability of the master.

If the XIV system determines that the RPO is greater than this value, the session state becomes Severe. You can specify an RPO between 30 and 86400 seconds. The default is 30 seconds.

##### **Synchronization schedule (HH:MM:SS)**

Select an interval for the creation of the XIV system synchronization schedule. The XIV system attempts to form consistent points of data by taking automatic snapshots of the volumes in the session at this interval. The default is **Minimum Interval**, which is 20 seconds.

If you select **Never**, synchronization is not scheduled and the XIV system does not create consistency groups. When the XIV system determines that the RPO threshold has been passed, the session state becomes Severe.

#### **ESS/DS Metro Global Mirror Options**

Select the Metro Global Mirror options that you want to associate with this session. These options are available only for TotalStorage Enterprise Storage Server Model 800, System Storage DS8000, or System Storage DS6000 storage systems in a Metro Global Mirror session:

##### **Consistency group interval time (seconds)**

Type how often, in seconds, the Metro Global Mirror session attempts to form a consistency group. When you reduce this value, it might be possible to reduce the data exposure of the session. A lower value causes the session to attempt to create consistency groups more frequently, which might affect the storage systems (for example, by increasing the processing load and message traffic load).

##### **Reset Secondary Reserve**

Select this option to remove any persistent reserves that might be set on the target volumes of the copy sets being started when a Start command is issued for the session.

##### **Note:**

Before enabling the Reset Secondary Reserves option, be aware that this action causes the session to overwrite all data on the target volume.

##### **Fail MM/GC if target is online (CKD only)**

Select this option to fail any session commands for a copy set if the session is online and if the target volume in the copy set is visible to a host.

## Metro Mirror Suspend Policy

You must select either **Hold I/O after Suspend** or **Release I/O after Suspend** when a replication link suspends. A replication link suspends for various reasons. In some cases, it happens because the replication link is down or the auxiliary storage controller has failed. In other cases, it might be the first failure in an eventual full failure at the primary site.

If your primary goal is zero data loss, specify **Hold I/O after Suspend**. Selecting this option means that no additional updates are made to storage until the entire scope of the failure can be determined. If you select this option, the systems by using that storage become unavailable until a decision is made about how to proceed. It might also require restarting the systems.

If your primary goal is availability, specify **Release I/O after Suspend**. This option enables workloads to continue running after all the replication links have been suspended, maintaining a consistent copy of data at the remote site. If the failure spreads through the primary site after the workloads resume running and results in a site failure, data loss occurs. The I/O requests that are issued after the resume and before the failure at the primary site are lost. The storage at the secondary site remains consistent, however, and you can recover it there.

### Hold I/O after Suspend

Select this option to block the application from writing while a consistent copy of the data is maintained on the remote site; however, it does not automatically release the application. This option keeps the source equal to the target. You must use the **Release I/O** command on the session or wait for the Hardware Freeze Timeout timer to expire before the application can continue to write to the source.

### Release I/O after Suspend

Select this option to block writes to the application while a consistent copy of the data is formed on the remote site. This operation is followed immediately by releasing the block so that the application can continue writing to the source. This option allows for little application impact, but causes the source to potentially be different from the target.

This option is the default setting for all new sessions.

## Manage H1-H2 with HyperSwap

Select this option to trigger a HyperSwap, redirecting application I/O to the secondary volumes, when there is a failure on the host accessible volumes. For Metro Global Mirror sessions, the Global Mirror portion of the session continues to run uninterrupted.

Tivoli Storage Productivity Center for Replication can manage the H1-H2 sequence of a Metro Mirror or Metro Global Mirror session by using HyperSwap if Tivoli Storage Productivity Center for Replication is running on z/OS, and volumes are attached by a Fibre Channel connection.

### Notes:

- When this option is selected, the Suspend H1-H2 command is available only if the Disable HyperSwap option is also selected.
- Setting this option automatically sets the Release IO after suspend Metro Mirror policy.

#### **Disable HyperSwap**

Select this option to prevent a HyperSwap from occurring by command or event.

#### **On Configuration Error:**

##### **Partition the system(s) out of the sysplex**

Select this option to partition out of the sysplex when a new system is added to the sysplex and encounters an error in loading the configuration. Restart the system if you select this option.

##### **Disable HyperSwap**

Select this option to prevent a HyperSwap from occurring by command or event.

#### **On Planned HyperSwap Error:**

##### **Partition out the failing system(s) and continue swap processing on the remaining system(s)**

Select this option to partition out the failing system and continue the swap processing on any remaining systems.

##### **Disable HyperSwap after attempting backout**

Select this option to stop the HyperSwap action, and disable the HyperSwap commands or events.

#### **On Unplanned HyperSwap Error:**

##### **Partition out the failing system(s) and continue swap processing on the remaining system(s)**

Select this option to partition out the failing systems and continue HyperSwap processing on the remaining systems when a new system is added to the sysplex and HyperSwap does not complete. You must restart the system if you select this option.

##### **Disable HyperSwap after attempting backout**

Select this option to stop the HyperSwap action, and disable the HyperSwap commands or events.

#### **ESS / DS FlashCopy Options for Role pair H3-I3**

The following option is available only for System Storage DS8000 version 4.2 or later.

##### **No Copy**

Select this option if you do not want the hardware to write the background copy until the source track is written to. Replication is done by using a copy-on-write technique. This technique does not copy data to the I3 volume until the blocks or tracks of the H3 volume are modified. The point-in-time volume image is composed of the unmodified data on the H3 volume, and the data that was copied to the I3 volume. If you want a complete point-in-time copy of the H3 volume to be created on the

I3 volume, do not use the No Copy option. This option causes the data to be asynchronously copied from the H3 volume to the I3 volume.

**Note:** Although you can select any space-efficient volume as the target, you cannot change the Permit Space Efficient Target flag. This flag is always set. When selecting space-efficient volumes as targets, you might receive an x0FBD error message if you attempt a full background copy. To avoid this message, select the No Copy option.

## Role pair status and progress

This topic describes the status messages for role pair status and progress.

Tivoli Storage Productivity Center for Replication provides detailed role pair status and progress messages for sessions. The status and progress messages are updated to provide a message that indicates what the session is currently doing at the time. By hovering over a progress bar, you can see specific information about the current action running on the session. Some status messages might include an estimated time to completion for the action in hours and minutes.

Role pair status is not provided for the XIV system Snapshot sessions because role pairs are not used for these sessions.

The status messages appear in the Session Details and Role Pair Details panel.

*Table 51. Detailed status messages for Participating and Non-Participating role pairs*

| Supported session type                                    | Status message                                                                              |
|-----------------------------------------------------------|---------------------------------------------------------------------------------------------|
| FlashCopy<br>Metro Mirror<br>Global Copy<br>Global Mirror | Starting <i>role_pair_name</i> relationships on the hardware                                |
| Metro Mirror<br>Global Mirror                             | Waiting for all pairs in the role pair <i>role_pair_name</i> to reach state of <i>state</i> |
| FlashCopy<br>Metro Mirror<br>Global Copy<br>Global Mirror | Terminating all pairs in role pair <i>role_pair_name</i>                                    |
| FlashCopy<br>Metro Mirror<br>Global Mirror                | Recovering all pairs in role pair <i>role_pair_name</i>                                     |
| Metro Mirror<br>Global Copy<br>Global Mirror              | Suspending all pairs in role pair <i>role_pair_name</i>                                     |
| FlashCopy                                                 | Background copy is running for role pair <i>role_pair_name</i>                              |
| Global Copy                                               | Waiting for all pairs in role pair <i>role_pair_name</i> to become consistent               |
| Global Copy                                               | Waiting for all pairs in role pairs <i>role_pair_name</i> to complete the initial copy      |
| FlashCopy                                                 | Waiting for all pairs in role pairs <i>role_pair_name</i> to complete FRR                   |

Table 51. Detailed status messages for Participating and Non-Participating role pairs (continued)

| Supported session type | Status message                                                                              |
|------------------------|---------------------------------------------------------------------------------------------|
| Global Mirror          | Waiting for all pairs in role pairs <i>role_pair_name</i> to join the Global Mirror session |

## Viewing session properties

This topic describes how to view session properties.

To view a session's properties, perform the following steps:

1. In the navigation tree, select **Sessions**. The Sessions panel is displayed.
2. Click the name of the session with the properties you want to view.
3. Select **View/Modify Properties** from the **Actions** menu and click **Go**.

## Viewing session details

You can view detailed information about a session, including role pairs, error count, whether the session is recoverable, copying progress, session type, and the timestamp.

Perform these steps to view session details:

1. In the navigation tree, select **Sessions**.
2. Select the session that you want to view.
3. Select **View Details** from the **Actions** drop-down menu, and click **Go**.

## Viewing additional details for Global Mirror sessions:

Additional detail information is available for Global Mirror sessions, including information about the Global Mirror master, consistency groups that have been formed, and data exposure.

Perform these steps to view additional details for Global Mirror sessions:

1. In the IBM Tivoli Storage Productivity Center for Replication navigation tree, select **Sessions**.
2. Select the Global Mirror session that you want to view.
3. Select **View Details** from the **Actions** drop-down menu, and click **Go**.
4. Click the **Global Mirror Info** tab. The following information is displayed on the tab:

### Global Mirror Master

Shows the name of the storage system acting as the Global Mirror master.

### Last Master Consistency Group Time

Shows the time that the last consistency group was formed

### Master Time During Last Query

Shows the time on the master storage device when the query was performed,

### Data Exposure

Shows the average exposure to potential data loss in seconds over the query interval.

**Session ID**

Shows the Global Mirror session ID.

**Master State**

Shows the state of the master session on the hardware.

**Unsuccessful CGs since last successful CG**

Shows the number of consistency groups that have failed to form since the last successful consistency group was formed.

**CG Interval Time**

Shows the interval time between attempts to form a consistency group.

**Max Coordination Interval**

Shows the extended distance consistency maximum coordination interval.

**Max CG Drain Time**

Shows the maximum time the consistent set of data is allowed to drain at the remote site before failing consistency group formation.

**Unsuccessful CGS/Previous Query**

Shows the number of consistency groups and percentage of consistency groups that were unsuccessful since the previous query.

**Unsuccessful CGS/Total**

Shows the total number of unsuccessful consistency groups and percentage of consistency groups that have failed.

**Successful CGS/Previous Query**

Shows the number of consistency groups and percentage of consistency groups that were successful since the previous query.

**Successful CGS/Total**

Shows the total number of successful consistency groups and percentage of consistency groups that have been successful.

**Consistency Group Failure Messages**

Shows the failure messages that have occurred on the Global Mirror session that prevented the formation of a consistency group.

**Data Exposure chart**

Shows the data exposure values in seconds for the last 15 minutes or 24 hours.

**Highlight Data Exposure**

Use the following fields to define a value in seconds for which you want data exposure to be tracked in the **Data Exposure** chart.

**Show Data Exposure over**

Data exposure that is above the value that is entered in this field is shown in the **Data Exposure** chart.

**Show Data Exposure under**

Data exposure that is below the value that is entered in this field is shown in the **Data Exposure** chart.

## Viewing storage system details

You can view detailed information about storage system, including the name, location, type, vendor, and the status of all connections to the storage system.

Perform these steps to view storage system details:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Perform one of these steps to view details for a specific storage system:
  - Click the storage system ID.
  - Select the storage system, click **View storage system details** from the **Actions** list, and then click **Go**.

## Viewing storage connection details

You can view storage connection details and a list of all storage systems that are located behind the connection.

Perform these steps to view storage connection details:

1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
2. Click the **Connections** tab.
3. Perform one of these steps to view details for a specific storage connection:
  - Click the storage connection ID.
  - Select the storage connection, click **View/modify Connection Details** from the actions list, and then click **Go**.

## Viewing volume details

You can view information about volumes such as the name of the volume, the capacity of the volume, and the type of volume.

1. In the navigation tree, select **Volumes**.
2. Select a storage system.
3. Depending on the type of storage, do one of the following:
  - a. Select **All IO Groups** or a specific I/O group.
  - b. Select **All Logical Storage Subsystems** or a specific logical storage subsystem.
  - c. Select **All Pools** or a specific pool.
4. Click **Perform Query**. Information about the volumes is displayed in a table.

## Viewing logical paths

You can view all logical paths that are defined on an IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, or IBM System Storage DS6000 storage system.

Perform one of these procedures to view logical paths:

- From the ESS/DS Paths panel of IBM Tivoli Storage Productivity Center for Replication:
  1. In the navigation tree, select **ESS/DS Paths**. The ESS/DS Paths panel is displayed.
  2. Click the storage system ID to display logical paths for that storage system.
- From the Storage Systems panel:
  1. In the navigation tree, select **Storage Systems**. The Storage Systems panel is displayed in the Storage Systems view.
  2. Select an ESS, DS6000, or DS8000 storage system for which you want to view logical paths.



3. Select **View Paths** from the **Select Action** list, and click **Go**. The ESS/DS Paths panel is displayed with a list of defined logical paths.

## Viewing console messages

This topic describes how to view the console and messages.

IBM Tivoli Storage Productivity Center for Replication provides detailed information about actions taken by users, errors that occur during normal operation, and hardware error indications.

From the graphical user interface, you can view console messages by selecting **Console** in the navigation tree. You can then click the link for the specific message code to get more information on the message.

You can also get detailed information and help for specific messages in the IBM Tivoli Storage Productivity Center information center at [publib.boulder.ibm.com/infocenter/tivihelp/v4r1/topic/com.ibm.tpc\\_V42.doc/welcome.html](http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/topic/com.ibm.tpc_V42.doc/welcome.html).

---

## Security

The IBM Tivoli Storage Productivity Center for Replication authentication process uses a configured user registry from either the operating system or Lightweight Directory Access Protocol (LDAP) server. To perform a specific action and manage specific sessions in the IBM Tivoli Storage Productivity Center for Replication GUI or CLI, the user must also have the appropriate authorization. Authorization is granted by assigning a specific role to the user account or user group.

## Users and groups

For authentication and authorization, IBM Tivoli Storage Productivity Center for Replication uses users and groups that are defined in a configured user registry on the management server, which is associated with either the local operating system or a Lightweight Directory Access Protocol (LDAP) server.

IBM Tivoli Storage Productivity Center for Replication does not provide the capability to create, update, or delete users or groups in the user registry. To manage users or groups, you must use the appropriate tool associated with the user registry in which the users and groups are stored.

IBM Tivoli Storage Productivity Center for Replication uses roles to authorize users to manage certain sessions and perform certain actions.

For more information about authentication, see information about single sign-on in the IBM Tivoli Storage Productivity Center documentation.

## Primary administrative ID

If you switch the authentication method, either from the local operating system to an LDAP server or vice versa, IBM Tivoli Storage Productivity Center for Replication removes all access to existing users and user groups. This occurs because the user IDs might not be on the same local operating system and the LDAP server; however, you must have at least one user ID that can log in to IBM Tivoli Storage Productivity Center for Replication.

When you change the authentication method using Tivoli Integrated Portal, you can specify a primary administrative ID for both local operating system and LDAP

authentication. Use this primary administrator to log in to IBM Tivoli Storage Productivity Center for Replication and manually add user IDs requiring access to IBM Tivoli Storage Productivity Center for Replication.

You can log in to both IBM Tivoli Storage Productivity Center for Replication and IBM Tivoli Storage Productivity Center using the primary administrative ID and password.

You cannot use the following characters for the IBM Tivoli Storage Productivity Center for Replication administrative password:

- square brackets ([ and ])
- semicolon (;)
- backward slash (\)

## User roles

A *user role* is a set of privileges that is assigned to a user or user group to allow the user or user group to perform certain tasks and manage certain sessions.

To be assigned to a role, each user or group of users must have a valid user ID or group ID in the user registry on the management server.

Both individual users and a group of users can be assigned to a role. All users in a group are assigned the role of the group. If a user is assigned to one role as an individual and a different role as a member of a group, the user has access to the permissions of the role that has greater access.

Restricting access to sessions prevents unwarranted administrative access. This is especially useful in an open environment, where there can be many storage administrators who are responsible for their servers, applications, databases, file systems, and so on.

IBM Tivoli Storage Productivity Center for Replication provides a set of predefined user roles: monitor, session operator, and administrator.

### Monitor

Monitors can view the health and status in the IBM Tivoli Storage Productivity Center for Replication GUI and CLI; however, they cannot modify or perform any commands or actions.

Monitors can view the following information:

- All storage systems and storage system details
- All connections and connection details
- All sessions and session details
- All path information
- Management server status and details

### Operator

Session operators can manage sessions to which they have been assigned, including:

- Adding or removing a session. The user ID that created the session is automatically granted access to manage that session.

- Performing actions on an assigned session, such as start, flash, terminate, and suspend.
- Modifying session properties.
- Adding copy sets to a session. The session operator can add volumes to a copy set only when the volume is not protected and not in another session.
- Removing copy sets from a session.
- Adding Peer To Peer Remote Copy (PPRC) paths, and removing paths with no hardware relationships. PPRC paths are a common resource used in IBM Tivoli Storage Productivity Center for Replication sessions and also in an ESS, DS6000, or DS8000 storage-system relationship that is established between two common logical subsystems (LSSs).

**Note:** The session operator cannot issue a force removal of a path.

**Note:** A path can also be auto-generated when starting a session.

- Monitoring health and status, including viewing the following information:
  - All storage systems and storage system details
  - All connections and connection details
  - All sessions and session details
  - All path information
  - Management server status and details

**Note:** Session operators can make changes only to the volumes that they own. They are not able to make changes to volumes being managed by other users.

## Administrator

During installation of IBM Tivoli Storage Productivity Center for Replication, the installation wizard requests an ID to use for the initial administrator user ID.

Administrators have unrestricted access. They can manage all sessions and perform all actions associated with IBM Tivoli Storage Productivity Center for Replication, including:

- Granting permissions to users and groups of users.
- Adding or removing a session. The user ID that created the session is automatically granted access manage that session.
- Performing actions on all sessions, such as start, flash, terminate, and suspend.
- Modifying session properties.
- Adding and removing copy sets from a session. The administrator can add volumes to a copy set only when the volume is not protected and not in another session.
- Protecting volumes and removing volume protection.
- Adding or removing storage system connections.
- Modifying connection properties.
- Assigning or changing storage system locations.
- Adding PPRC paths and removing paths with no hardware relationships. PPRC paths are a common resource used in IBM Tivoli Storage Productivity Center for Replication sessions and also in an ESS, DS6000, or DS8000 storage-system relationship that is established between two common logical subsystems (LSSs).

**Note:** A path can also be auto-generated when starting a session.

- Managing management servers. The standby management server is a common resource that is available to multiple sessions.
- Packaging program error (PE) log files.
- Monitoring health and status, including viewing the following information:
  - All storage systems and storage system details
  - All connections and connection details
  - All sessions and session details
  - All path information
  - Management server status and details

**Important:** IBM Tivoli Storage Productivity Center supports multiple user roles, including the Superuser role. A superuser can perform all IBM Tivoli Storage Productivity Center functions. For IBM Tivoli Storage Productivity Center superusers to have full access to IBM Tivoli Storage Productivity Center for Replication, the Superuser group must be added to the IBM Tivoli Storage Productivity Center for Replication and assigned to the Administrator role. Then, you can manage the IBM Tivoli Storage Productivity Center and IBM Tivoli Storage Productivity Center for Replication products by groups, instead of by user IDs.

**Note:** Administrators cannot revoke their own administrative access rights.

## Adding the IBM Tivoli Storage Productivity Center for Replication Administrator role to the IBM Tivoli Storage Productivity Center Superuser group

If you use local operating system or Lightweight Directory Access Protocol (LDAP) authentication, you must add the IBM Tivoli Storage Productivity Center superuser group to IBM Tivoli Storage Productivity Center for Replication with administrator privileges.

**Prerequisite:** You must have Administrator privileges to perform this action.

Perform these steps to add the superuser group to IBM Tivoli Storage Productivity Center for Replication:

1. Log in to IBM Tivoli Storage Productivity Center as the superuser.
2. In the navigation tree, expand **Administrative Services > Configuration** and select **Role-to-Group Mappings**.
3. Locate the name of the superuser user group.
4. Log in to IBM Tivoli Storage Productivity Center for Replication using a user ID with administrator privileges.
5. In the navigation tree, select **Administration**. The Administration panel is displayed.
6. Click **Add Access**. The Add Access wizard is displayed.
7. Type the name of the superuser group in the **User or group names** field, and click **Next**.
8. Select the name of the superuser group, and click **Next**.
9. Select **Administrator privileges**, and click **Next**.
10. Click **Next** to confirm this action.
11. Click **Finish**.

## Granting access privileges for a user

You can assign user roles to an IBM Tivoli Storage Productivity Center for Replication user to grant access privileges to individual sessions and tasks.

Perform the following steps to authorize a user:

1. Create the user ID or group ID if it does not already exist in the user registry, either the operating system of the active management server or Lightweight Directory Access Protocol (LDAP) server.
2. Log in to IBM Tivoli Storage Productivity Center for Replication as a user with administrator privileges.
3. In the navigation tree, select **Administration**. The Administration panel is displayed.
4. Click **Add Access**. The Add Access wizard is displayed.
5. Type the name of the user to whom you want to give access, and click **Next**. The Select Users and Groups panel is displayed.

**Tip:** You can enter a partial name and use the \* wildcard character to represent zero or more characters.

6. Select one or more names from the list of found users.
7. Select the role to associate with this user.
8. If you selected the Operator role, select one or more session that this user can manage, and click **Next**.
9. Click **Next** to confirm this action.
10. Click **Finish**.

## Viewing access privileges for a user

You can view a list of all IBM Tivoli Storage Productivity Center for Replication user and their assigned roles. You can also view the assigned sessions for each user.

Perform the following steps to view access privileges for a user:

1. Log in to IBM Tivoli Storage Productivity Center for Replication as a user with administrator privileges.
2. In the navigation tree, select **Administration**. The Administration page is displayed with a list of IBM Tivoli Storage Productivity Center for Replication users and user groups and their associated role.
3. Select the user whose access privileges you want to view.
4. Select **View/Modify Access** from the **Actions** drop-down list, and click **Go**. The View/Modify Access panel is displayed. This panel shows the role assigned to the user and lists the sessions that the user can manage.
5. Click **Cancel**.

## Modifying access privileges for a user

You can change the user role and assigned sessions for an IBM Tivoli Storage Productivity Center for Replication user.

**Prerequisite:** You must have Administrator privileges to perform this action.

Perform the following steps to modify the access privileges for a user:

1. Log in to IBM Tivoli Storage Productivity Center for Replication as a user with administrator privileges.
2. In the navigation tree, select **Administration**. The Administration panel is displayed with a list of users and user groups and their associated role
3. Select the user whose access privileges you want to view.
4. Select **View/Modify Access** from the **Select Action** drop-down list, and click **Go**. The View/Modify Access panel is displayed. This panel shows the role assigned to the user and lists the sessions that the user can manage.
5. Select the role to associate with this user.
6. If you selected the Operator role, select one or more session that this user can manage and click **Next**.
7. Click **OK**.

## Removing access privileges for a user

You can remove access privileges for an IBM Tivoli Storage Productivity Center for Replication user. When you remove access, the user ID cannot access the IBM Tivoli Storage Productivity Center for Replication GUI or run commands from the command line.

**Prerequisite:** You must have Administrator privileges to perform this action.

Perform the following steps to remove user access:

1. Log in to IBM Tivoli Storage Productivity Center for Replication as a user with administrator privileges.
2. In the navigation tree, select **Administration**. The Administration panel is displayed with a list of users and user groups and their associated role.
3. Select the user from which you want to remove access.
4. Select **Remove Access** from the **Actions** list, and click **Go**.

---

## Chapter 6. Reporting

Use Tivoli Storage Productivity Center reports to view detailed information about the storage resources in your environment. These reports are organized into different types and categories and provide both summary and detailed information depending on your needs.

Data in reports is based on statistical information accumulated by data collection schedules and stored in the database repository. You can select the specific systems, storage subsystems, switches, fabrics, NAS filers, clusters, disks, file systems, and even users upon which you want to report.

Reports are organized into report types and categories that you can expand to reveal additional subcategories of reporting. Many reports are also accessible through the topology viewer.

- **Reporting types:** system reports, batch reports, user ID reports, user-defined online reports
- **Report categories:** asset, availability, capacity, usage, usage violation, backup, storage subsystem performance, switch performance

Consider the following steps when determining how to best use Tivoli Storage Productivity Center reports:

1. Determine how you want to generate and view reports. See “Choosing a reporting type” for more information.
2. Identify which reports contain the most important information about your storage environment. See “Choosing a report category” on page 771 for more information.

---

### Choosing a reporting type

Use pre-defined (system reports), saved reports, batch reports, rollup reports, and user-defined reports, tables, and charts to view the information collected by monitoring and quota jobs. Tivoli Storage Productivity Center includes these different reporting types to provide you with the flexibility to view data about your storage resources according to the needs of your environment.

The following table describes the reporting types available in Tivoli Storage Productivity Center. Use this table to help determine how to best deliver and view reports about your storage resources.

| Reporting type                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Location in the navigation tree                                                                                                                                                                                                 |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>User-defined online reports</b> | <p>These reports can be generated at any time during a product session and are only available for that session. Each manager within Tivoli Storage Productivity Center has a set of reports specific to that manager. For example, to view detailed reports about storage subsystems, access the reporting function within Disk Manager.</p> <p>See the following topics for more information about online reports:</p> <ul style="list-style-type: none"> <li>• “User-defined online reports” on page 793</li> <li>• “Data Manager reports” on page 804</li> <li>• “Using Data Manager for Databases reports” on page 933</li> <li>• “Disk Manager reports” on page 961</li> <li>• “Fabric Manager reports” on page 973</li> </ul>                                           | <p>Data Manager &gt; Reporting</p> <p>Data Manager for Databases &gt; Reporting</p> <p>Disk Manager &gt; Reporting</p> <p>Fabric Manager &gt; Reporting</p>                                                                     |
| <b>System reports</b>              | <p>System reports are automatically generated by Tivoli Storage Productivity Center and contain some of the most common information about your monitored storage resources. These pre-defined reports enable you to view your storage enterprise data quickly and efficiently with one click. Data for these system reports are gathered every time that monitoring jobs are run against your storage resources.</p> <p>See the following topics for more information about online reports:</p> <ul style="list-style-type: none"> <li>• “System reports - Data” on page 981</li> <li>• “System reports - Data Manager for Databases” on page 993</li> <li>• “System Reports - Disk Manager” on page 999</li> <li>• “System reports - Fabric Manager” on page 1016</li> </ul> | <p>IBM Tivoli Storage Productivity Center &gt; Reporting &gt; System Reports</p>                                                                                                                                                |
| <b>User ID (saved) reports</b>     | <p>Create saved reports when you want specific reports generated for you whenever you log in to a product session. This enables you to save reports that you can generate at any time without having to define your report criteria again. Saved reports are only visible to the user ID that created them.</p> <p>See “User ID (saved) reports” on page 1034 for more information.</p>                                                                                                                                                                                                                                                                                                                                                                                       | <p>IBM Tivoli Storage Productivity Center &gt; Reporting &gt; My Reports &gt; <i>user ID</i> Reports</p> <p>where <i>user ID</i> represents the IBM Tivoli Storage Productivity Center ID of the user that saved the report</p> |



| Reporting type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Location in the navigation tree                                                 |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------|
| <b>Batch reports</b>       | <p>A batch report represents any Tivoli Storage Productivity Center report that you want to run on a regular basis and save its data to a file. You can view the information in a generated batch report file directly or import it into an application of your choice.</p> <p>To create a batch report, you must determine which Tivoli Storage Productivity Center report you want to generate, the schedule on which to generate the report, and the file to which to save its data.</p> <p>You can save the data generated for a batch report to any of the following file formats: CSV, HTML, formatted text, PDF chart, and HTML chart.</p> <p>See “Batch reports” on page 1021 for more information.</p>                         | IBM Tivoli Storage Productivity Center > Reporting > My Reports > Batch Reports |
| <b>Rollup Reports</b>      | <p>Use rollup reports to view the storage information collected by the master Tivoli Storage Productivity Center server from its associated subordinate servers. Rollup reports enable you to have a network-wide perspective of storage usage in an environment where multiple Tivoli Storage Productivity Center servers are deployed and collecting storage metrics.</p> <p>See the following topics for more information about online reports:</p> <ul style="list-style-type: none"> <li>• “Rollup reports” on page 1035</li> <li>• “Tivoli Storage Productivity Center servers” on page 278</li> <li>• “Collecting storage statistics from Tivoli Storage Productivity Center servers (TPC server probes)” on page 340</li> </ul> | IBM Tivoli Storage Productivity Center > Reporting > Rollup Reports             |
| <b>Data Source reports</b> | <p>Use these reports to view detailed information about the agents from which Tivoli Storage Productivity Center collects information about storage resources. Data sources include CIMOM agents, Storage Resource agents, Data agents, Fabric agents, VMWare, and subordinate Tivoli Storage Productivity Center servers.</p> <p>See Data Source reports for more information.</p>                                                                                                                                                                                                                                                                                                                                                     | IBM Tivoli Storage Productivity Center > Reporting > Data Source Reports        |

## Choosing a report category

IBM Tivoli Storage Productivity Center provides a number of categories for reports that help you identify which reports contain the information you want to view about your storage environment. Each category of report contains information gathered by different data collection and quota jobs.

The following table describes the reporting categories. Use this table to help determine which reports are most useful to you and which data collection jobs to run for those reports.

| Report Category                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | What data collection job should I run for this report?                                                            |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------|
| <p><b>Asset reports</b></p>        | <p>View itemized statistics about your storage resources in a hierarchical view that enables you to drill down to view those resources in greater detail.</p> <p><b>Location:</b> These reports are available in the following locations on the navigation tree:</p> <ul style="list-style-type: none"> <li>• Data Manager &gt; Reporting &gt; Asset</li> <li>• Data Manager for Databases &gt; Reporting &gt; Asset</li> <li>• IBM Tivoli Storage Productivity Center &gt; Reporting &gt; My Reports (if you save the definition of an online report it appears in this location)</li> <li>• Data Manager for Databases &gt; My Reports</li> </ul> <p><b>Example:</b> View information about a storage subsystem and drill down into its related storage pools, disks, volumes, and disk groups.</p> <p>See “Asset reports” on page 807 for more information about Asset reports.</p> | <p>Probes, Discovery</p> <p>Data for System-wide Users and Monitored Directory reports are gathered by Scans.</p> |
| <p><b>Availability reports</b></p> | <p>View availability statistics about your storage resources, both from a network point of view and from a computer uptime perspective.</p> <p><b>Location:</b> These reports are available in the following locations on the navigation tree:</p> <ul style="list-style-type: none"> <li>• Data Manager &gt; Reporting &gt; Availability</li> <li>• IBM Tivoli Storage Productivity Center &gt; Reporting &gt; My Reports (if you save the definition of an online report it appears in this location)</li> </ul> <p><b>Example:</b> Use a Computer Uptime report to determine the percentage of time that a computer was up, the number of times it was rebooted, and the total time it was down within a user-defined period of time.</p> <p>See “Availability Reports” on page 831 for more information about Availability reports.</p>                                            | <p>Ping data: Pings</p> <p>Computer Uptime: Probes</p>                                                            |

| Report Category               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | What data collection job should I run for this report? |
|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------|
| <p>TPC-wide Storage Space</p> | <p>Examine storage capacities of storage resources at a file system, computer, or network level.</p> <p><b>Location:</b> These reports are available in the following locations on the navigation tree:</p> <ul style="list-style-type: none"> <li>• Data Manager &gt; Reporting &gt; TPC-wide Storage Space</li> <li>• Data Manager for Databases&gt; Reporting &gt; Capacity</li> <li>• IBM Tivoli Storage Productivity Center &gt; Reporting &gt; My Reports (if you save the definition of an online report it appears in this location)</li> <li>• Data Manager for Databases &gt; My Reports</li> </ul> <p><b>Example:</b> Identify how much storage space is being used on the file systems within your environment, and how much unused storage space is available.</p> <p>See “TPC-wide Storage Space Reports” on page 832 for more information about Capacity reports.</p>                                                                                                                                     | <p>Probes</p>                                          |
| <p>Usage reports</p>          | <p>View information about the usage and growth of consumption of your storage resources, as well as select files against which to run an Tivoli Storage Productivity Center archive/backup job.</p> <p><b>Location:</b> These reports are available in the following locations on the navigation tree:</p> <ul style="list-style-type: none"> <li>• Data Manager &gt; Reporting &gt; Usage</li> <li>• Data Manager for Databases &gt; Reporting &gt; Usage</li> <li>• IBM Tivoli Storage Productivity Center &gt; Reporting &gt; My Reports (if you save the definition of an online report it appears in this location)</li> <li>• Data Manager for Databases &gt; My Reports (if you save the definition of an online report it appears in this location)</li> </ul> <p><b>Example:</b> Determine how much of the storage being used in your environment can be considered wasted space (orphaned, obsolete, or forbidden files).</p> <p>See “Usage Reports” on page 892 for more information about Usage reports.</p> | <p>Scans</p>                                           |

| Report Category         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | What data collection job should I run for this report? |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------|
| Usage Violation reports | <p>View quota and constraint violations as well as select files for archive/backup jobs.</p> <p><b>Location:</b> These reports are available in the following locations on the navigation tree:</p> <ul style="list-style-type: none"> <li>• Data Manager &gt; Reporting &gt; Usage Violations</li> <li>• Data Manager for Databases &gt; Reporting &gt; Usage Violations</li> <li>• IBM Tivoli Storage Productivity Center &gt; Reporting &gt; My Reports (if you save the definition of an online report it appears in this location)</li> <li>• Data Manager for Databases &gt; My Reports (if you save the definition of an online report it appears in this location)</li> </ul> <p><b>Example:</b> Identify the users or groups of users that are using storage space on machines where they should not.</p> <p>See “Usage Violation reports” on page 899 for more information about Usage Violation reports.</p>                                                                           | Scans, Quotas, Constraints                             |
| Backup reports          | <p>Investigate the sizes of data that would be required to backup to achieve full backups, determine the amount of space required to perform an incremental backup of any or all of your data at any given point in time, and select files for archive/backup jobs.</p> <p><b>Location:</b> These reports are available in the following locations on the navigation tree:</p> <ul style="list-style-type: none"> <li>• Data Manager &gt; Reporting &gt; Backup</li> <li>• IBM Tivoli Storage Productivity Center &gt; Reporting &gt; My Reports (if you save the definition of an online report it appears in this location)</li> <li>• Data Manager for Databases &gt; My Reports (if you save the definition of an online report it appears in this location)</li> </ul> <p><b>Example:</b> Determine if you have enough storage space to perform a backup of your data (incremental or full).</p> <p>See “Backup reports” on page 901 for more information about Usage Violation reports.</p> | Scans                                                  |

| Report Category                                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | What data collection job should I run for this report? |
|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------|
| <p><b>Monitored Computer Storage Space</b></p>      | <p>Use these host-based reports to view the data that is collected by probes for computers, Network Attached Storage, and hypervisors. These reports contain detailed information about the consumption of file system space and enable you to relate that space back to external or internal storage without counting shared space multiple times.</p> <p><b>Location:</b> These reports are available in the following locations on the navigation tree:</p> <ul style="list-style-type: none"> <li>• Data Manager &gt; Reporting &gt; Monitored Computer Storage Space</li> <li>• IBM Tivoli Storage Productivity Center &gt; Reporting &gt; My Reports (if you save the definition of an online report it appears in this location)</li> </ul> <p>See “Monitored Computer Storage Space reports” on page 903 for more information about Monitored Computer Storage Space reports.</p> | <p>Probes</p>                                          |
| <p><b>Storage subsystem reports</b></p>             | <p>View storage capacity for a computer, file system, storage subsystem, volumes, array sites, ranks, and disks. These reports also enable you to view the relationships among the components of a storage subsystem.</p> <p><b>Location:</b> Disk Manager &gt; Reporting &gt; Storage Subsystems</p> <p><b>Example:</b> Determine if you have enough storage space to perform a backup of your data (incremental or full).</p> <p>See “Disk Manager reports” on page 961 for more information about Storage Subsystem reports.</p>                                                                                                                                                                                                                                                                                                                                                       | <p>Probes</p>                                          |
| <p><b>Storage Subsystem Performance reports</b></p> | <p>View performance data for storage subsystems.</p> <p><b>Location:</b> Disk Manager &gt; Reporting &gt; Storage Subsystem Performance</p> <p><b>Example:</b> Check the performance rate for storage subsystem activities, such as read I/O rate, read data rate, total port I/O rate, etc.</p> <p>See “Disk Manager reports” on page 961 for more information about Storage Subsystem Performance Monitor reports.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <p>Subsystem Performance Monitors</p>                  |

| Report Category                                                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | What data collection job should I run for this report?                                                          |
|--------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| Switch Performance reports                                         | <p>View performance data for fabric switches.</p> <p><b>Location:</b> Fabric Manager &gt; Reporting &gt; Switch Performance</p> <p><b>Example:</b> Check the performance rate for switch activities, such as port send packet rate, port send data rate, total port data rate, etc.</p> <p>See “Fabric Manager reports” on page 973 for more information about Switch Performance Monitor reports.</p>                                                                                                                                                                                                                                                                                                   | Switch Performance Monitors                                                                                     |
| Rollup reports: Asset, Capacity, Database Asset, Database Capacity | <p>View the storage information that has been collected by IBM Tivoli Storage Productivity Center servers through a single interface. These rollup reports enable you to have a consolidated perspective of storage usage in an environment where multiple IBM Tivoli Storage Productivity Center servers are deployed and collecting storage metrics.</p> <p><b>Location:</b> IBM Tivoli Storage Productivity Center &gt; Reporting &gt; Rollup Reports</p> <p><b>Example:</b> Check the disk capacity of all the computers monitored by different servers in your environment.</p> <p>See “Rollup reports” on page 1035 for more information about rollup reports.</p>                                 | IBM Tivoli Storage Productivity Center Server Probes                                                            |
| Data Source reports                                                | <p>Use data source reports to view detailed information about the agents from which Tivoli Storage Productivity Center collects information about storage resources. Data sources include CIMOM agents, Storage Resource agents, Data agents, Fabric agents, VMWare, and subordinate Tivoli Storage Productivity Center servers.</p> <p><b>Location:</b> IBM Tivoli Storage Productivity Center &gt; Reporting &gt; Data Source Reports</p> <p><b>Example:</b> View detailed information about the CIMOM Agents that have been added to Tivoli Storage Productivity Center as data sources.</p> <p>See “Data source reports - overview” on page 1093 for more information about data source reports.</p> | These reports are available for data sources that have been associated with Tivoli Storage Productivity Center. |

## What can I find out using reports?

Learn about some of the key storage questions that you can answer by generating Tivoli Storage Productivity Center reports.

Use the following table to view a list of reports that can be generated to answer key questions about storage resources that are monitored by Tivoli Storage Productivity Center.

| What do you want to know?                                                                                              | Which report answers this question?                                                                             | Which schedule can I run to generate this report? |
|------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|---------------------------------------------------|
| What hardware is currently being used within the enterprise?                                                           | Data Manager > Reporting > Asset > By <storage_resource_type>                                                   | Probes                                            |
| How much storage space is being used?                                                                                  | Data Manager > Reporting > TPC-wide Storage Space > Consumed File System Space                                  | Probes                                            |
| Of the space that is being used, how much is wasted space? (orphaned, obsolete, or duplicate files)                    | Data Manager > Reporting > Usage > Files<br>Data Manager > Reporting > Usage Violations > Constraint Violations | Scans                                             |
| How much unused or unallocated storage space is available?                                                             | Data Manager > Reporting > TPC-wide Storage Space > Available File System Space                                 | Probes                                            |
| Is there enough space to do a backup of our data? (incremental or full)                                                | Data Manager > Reporting > Backup > Backup Storage Requirements                                                 | Scans                                             |
| Which files in the environment are at risk if a disaster hits?                                                         | Data Manager > Reporting > Backup > Most at Risk Files                                                          | Scans                                             |
| What "forbidden files" are taking up storage space on the network?                                                     | Data Manager > Reporting > Usage Violations > Constraint Violations                                             | Scans                                             |
| Which user or group of users is consuming space on computers that they should not?                                     | Data Manager > Reporting > Usage Violations > Quota Violations                                                  | Scans                                             |
| How fast are storage resources being consumed?                                                                         | Data Manager > Reporting > TPC-wide Storage Space > Consumed File System Space                                  | Probes                                            |
| Which user or business unit is consuming the most storage resources?                                                   | Data Manager > Reporting > Usage > Users<br>Data Manager > Reporting > Usage > OS User Groups                   | Scans                                             |
| Where is disk space reaching a critical shortage?                                                                      | Data Manager > Reporting > TPC-wide Storage Space > Disk Space                                                  | Probes                                            |
| Which devices have been experiencing the most use/traffic? Which devices have been experiencing the least use/traffic? | Data Manager > Reporting > Usage > Access Load                                                                  | Scans                                             |
| Based on historical consumption, how much storage space should be added and when should it be added?                   | Data Manager > Reporting > TPC-wide Storage Space > Available File System Space                                 | Probes                                            |
| What files have not been used within a certain time period?                                                            | Data Manager > Reporting > Usage > Files > Most Obsolete Files                                                  | Scans                                             |
| What are the largest files or directories within the environment?                                                      | Data Manager > Reporting > Usage > Files > Largest Files                                                        | Scans                                             |

| What do you want to know?                                                                                                      | Which report answers this question?                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Which schedule can I run to generate this report? |
|--------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------|
| What are the most modified files within the environment?                                                                       | Data Manager > Reporting > Usage > Access Load > Modification Time                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Scans                                             |
| Which computers are “up” or “down” within the environment?                                                                     | Data Manager > Reporting > Availability > Ping                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Pings                                             |
| What hardware is most prone to error? (for example, disk defects)                                                              | Data Manager > Reporting > Asset > System-wide > Disks                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Probes                                            |
| Which file owners no longer exist in the environment?                                                                          | Data Manager > Reporting > Usage > Files > Orphan files                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Scans                                             |
| How much storage space is an instance consuming?                                                                               | <ul style="list-style-type: none"> <li>• <b>Oracle:</b> Data Manager for Databases &gt; Reporting &gt; Capacity &gt; Oracle &gt; Total Instance Storage</li> <li>• <b>SQL/Server:</b> Data Manager for Databases &gt; Reporting &gt; Capacity &gt; SQL/Server &gt; Total Instance Storage</li> <li>• <b>Sybase:</b> Data Manager for Databases &gt; Reporting &gt; Capacity &gt; Sybase &gt; Total Instance Storage</li> <li>• <b>UDB:</b> Data Manager for Databases &gt; Reporting &gt; Capacity &gt; UDB &gt; Total Instance Storage</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Probes                                            |
| Which user or business unit is consuming the most storage resources and on which instances are they consuming those resources? | <p><b>Oracle:</b></p> <ul style="list-style-type: none"> <li>• Data Manager for Databases &gt; Reporting &gt; Usage &gt; Oracle &gt; Users &gt; Group by User Name, <i>or</i></li> <li>• Data Manager for Databases &gt; My Reports &gt; System Reports &gt; Oracle - User Database Space Usage</li> </ul> <p><b>SQL/Server:</b></p> <ul style="list-style-type: none"> <li>• Data Manager for Databases &gt; Reporting &gt; Usage &gt; SQL/Server &gt; Users &gt; Group by User Name, <i>or</i></li> <li>• Data Manager for Databases &gt; My Reports &gt; System Reports &gt; SqlServer - User Database Space Usage</li> </ul> <p><b>Sybase:</b></p> <ul style="list-style-type: none"> <li>• Data Manager for Databases &gt; Reporting &gt; Usage &gt; Sybase &gt; Users &gt; Group by User Name, <i>or</i></li> <li>• Data Manager for Databases &gt; My Reports &gt; System Reports &gt; Sybase - User Database Space Usage</li> </ul> <p><b>UDB:</b></p> <ul style="list-style-type: none"> <li>• Data Manager for Databases &gt; Reporting &gt; Usage &gt; UDB &gt; Users &gt; Group By User Name, <i>or</i></li> <li>• Data Manager for Databases &gt; My Reports &gt; System Reports &gt; UDB - User Database Space Usage</li> </ul> | Scans                                             |



| What do you want to know?                                                                                                          | Which report answers this question?                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Which schedule can I run to generate this report? |
|------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------|
| Which tablespaces are reaching a critical shortage of free space?                                                                  | <ul style="list-style-type: none"> <li>• <b>Oracle:</b> Data Manager for Databases &gt; Reporting &gt; Capacity &gt; Oracle &gt; Data File Free Space &gt; By Tablespace</li> <li>• <b>UDB:</b> Data Manager for Databases &gt; Reporting &gt; Capacity &gt; UDB &gt; Container Free Space &gt; By Tablespace</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Probes                                            |
| Which tablespaces contain the most free space?                                                                                     | <ul style="list-style-type: none"> <li>• <b>Oracle:</b> Data Manager for Databases &gt; Reporting &gt; Capacity &gt; Oracle &gt; Data File Used Space &gt; By Tablespace</li> <li>• <b>UDB:</b> Data Manager for Databases &gt; Reporting &gt; Capacity &gt; UDB &gt; Container Used Space &gt; By Tablespace</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Probes                                            |
| Which databases are reaching a critical shortage of free space?                                                                    | <ul style="list-style-type: none"> <li>• <b>SQL/Server:</b> Data Manager for Databases &gt; Reporting &gt; Capacity &gt; SQL/Server &gt; Data File Free Space &gt; By Database</li> <li>• <b>Sybase:</b> Data Manager for Databases &gt; Reporting &gt; Capacity &gt; Sybase &gt; Fragment Free Space &gt; By Database</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Probes                                            |
| Which databases contain the most used space?                                                                                       | <ul style="list-style-type: none"> <li>• <b>SQL/Server:</b> Data Manager for Databases &gt; Reporting &gt; Capacity &gt; SQL/Server &gt; Data File Used Space &gt; By Database</li> <li>• <b>Sybase:</b> Data Manager for Databases &gt; Reporting &gt; Capacity &gt; Sybase &gt; Fragment Used Space &gt; By Database</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Probes                                            |
| Based on historical consumption of storage space for an RDBMS, how much storage space should be added and when should it be added? | <ul style="list-style-type: none"> <li>• <b>Oracle:</b> Data Manager for Databases &gt; Reporting &gt; Capacity &gt; Oracle &gt; Data File Free Space &gt; By Database &gt; History Chart:Free Space % for selected</li> <li>• <b>SQL/Server:</b> Data Manager for Databases &gt; Reporting &gt; Capacity &gt; SQL/Server &gt; Data File Free Space &gt; By Instance &gt; History Chart:Free Space % for selected</li> <li>• <b>Sybase:</b> Data Manager for Databases &gt; Reporting &gt; Capacity &gt; Sybase &gt; Device Free Space &gt; By Device &gt; History Chart:Free Space % for selected</li> <li>• <b>UDB:</b> Data Manager for Databases &gt; Reporting &gt; Capacity &gt; UDB &gt; Container Free Space &gt; By Instance &gt; History Chart:Free Space % for selected</li> </ul> | Probes                                            |

| What do you want to know?                                                                                      | Which report answers this question?                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Which schedule can I run to generate this report? |
|----------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------|
| <p>What are the largest tables within the environment?</p>                                                     | <ul style="list-style-type: none"> <li>• <b>Oracle:</b> Data Manager for Databases &gt; Reporting &gt; Usage &gt; Oracle &gt; Tables &gt; Largest Tables</li> <li>• <b>SQL/Server:</b> Data Manager for Databases &gt; Reporting &gt; Usage &gt; SQL/Server &gt; Tables &gt; Largest Tables</li> <li>• <b>Sybase:</b> Data Manager for Databases &gt; Reporting &gt; Usage &gt; Sybase &gt; Tables &gt; Largest Tables</li> <li>• <b>UDB:</b> Data Manager for Databases &gt; Reporting &gt; Usage &gt; UDB &gt; Tables &gt; Largest Tables</li> </ul> | Scans                                             |
| <p>Which segments in an instance contain the most extents?</p>                                                 | <p><b>Oracle:</b> Data Manager for Databases &gt; Reporting &gt; Usage &gt; Oracle &gt; Segments &gt; Most Extents</p>                                                                                                                                                                                                                                                                                                                                                                                                                                 | Scans                                             |
| <p>How much allocated space for an object is being wasted?</p>                                                 | <p><b>Oracle:</b> Data Manager for Databases &gt; My Reports &gt; System Reports &gt; Oracle - Segments with Wasted Space</p>                                                                                                                                                                                                                                                                                                                                                                                                                          | Scans                                             |
| <p>How many free extents does a tablespace have?</p>                                                           | <p><b>Oracle:</b> Data Manager for Databases &gt; Reporting &gt; Asset &gt; By DBMS Type &gt; Oracle &gt; instance_name &gt; By Tablespace</p>                                                                                                                                                                                                                                                                                                                                                                                                         | Probes                                            |
| <p>How many tables, indexes, and datafiles/containers are associated with a tablespace?</p>                    | <ul style="list-style-type: none"> <li>• <b>Oracle:</b> Data Manager for Databases &gt; Reporting &gt; Asset &gt; System-wide &gt; Oracle &gt; Tablespaces</li> <li>• <b>UDB:</b> Data Manager for Databases &gt; Reporting &gt; Asset &gt; System-wide &gt; UDB &gt; Tablespaces</li> </ul>                                                                                                                                                                                                                                                           | Scans                                             |
| <p>How many tables, indexes, and datafiles/fragments are associated with a database?</p>                       | <ul style="list-style-type: none"> <li>• <b>SQL/Server:</b> Data Manager for Databases &gt; Reporting &gt; Asset &gt; System-wide &gt; SQL/Server &gt; Databases</li> <li>• <b>Sybase:</b> Data Manager for Databases &gt; Reporting &gt; Asset &gt; System-wide &gt; Sybase &gt; Databases</li> </ul>                                                                                                                                                                                                                                                 | Scans                                             |
| <p>How can I determine which users have violated a quota on a specific instance?</p>                           | <p><b>All DBMSs:</b> Data Manager for Databases &gt; Reporting &gt; Usage Violations &gt; Quota Violations &gt; Instance Quotas &gt; By Instance</p>                                                                                                                                                                                                                                                                                                                                                                                                   | Scans                                             |
| <p>What are the sizes and locations of Control Files and log files for an instance within the environment?</p> | <ul style="list-style-type: none"> <li>• <b>Oracle:</b> Data Manager for Databases &gt; Reporting &gt; Asset &gt; By DBMS Type &gt; Oracle &gt; instance_name</li> <li>• <b>SQL/Server:</b> Data Manager for Databases &gt; Reporting &gt; Asset &gt; By DBMS Type &gt; SQL/Server &gt; instance_name &gt; Databases</li> </ul>                                                                                                                                                                                                                        | Probes                                            |
| <p>How can I determine how close a user is to exceeding a quota?</p>                                           | <p><b>All DBMSs:</b> Data Manager for Databases &gt; Usage Violations &gt; Quota Violations &gt; All Quotas &gt; By User</p>                                                                                                                                                                                                                                                                                                                                                                                                                           | Scans                                             |

---

## Report considerations and limitations

Keep in mind these considerations and limitations when viewing reports in Tivoli Storage Productivity Center.

### Data collection for reports

Some columns within reports might not be populated with data for the storage resources that you are monitoring. Columns appear empty if you have not configured Tivoli Storage Productivity Center to discover or collect the information intended for that column. Use discovery, probe, ping, or scan schedules to ensure that you collect information for all the columns that appear in reports.

When you specify a start time for using performance monitor data in a report, the start time might not coincide with the start of a performance monitor data collection interval. If the requested start time is after the start of a data collection interval and the duration of the data collection interval is long enough so that it ends at or after the requested start time, Tivoli Storage Productivity Center includes all data from the collection interval that began before the requested start time. This might include data collected before the start time that you specified. This processing ensures that the maximum amount of available performance data is utilized. For example, if a performance monitor collects data every 20 minutes starting at 01:00, but you request performance data starting at 1:15, Tivoli Storage Productivity Center still analyzes performance data starting at 01:00 that has a duration of 20 minutes.

### Report charts that contain long path names

If you generate charts for reports and these charts contain long path names or a large number of segments, the text on the charts might overlap and make them difficult to read.

### Information about directories

Before information about directories is displayed in reports, you must complete the following tasks:

1. Include directories that have been discovered by Tivoli Storage Productivity Center in a directory group.
2. Include the directory group in a data collection scan.

After a scan schedule has collected data about the directories in a directory group, information about the directories is displayed in By Directory reports

### Storage totals for Network-wide reports

Tivoli Storage Productivity Center provides a number of network-wide reports that display storage totals based on the data collected from the monitored computers and storage subsystems in your network. Click the magnifying glass icon next to a row in a network-wide report to view the details about the computers that are included in those storage totals. However, the reports that you drill down into display information about computers only. Storage information about storage subsystems is not included in the values for reports that are shown when you drill down from a network-wide report.

For example, the Data Manager > Reporting > Capacity > Disk Capacity > Network-wide report displays disk capacity totals for the computers and storage subsystems in your network that are monitored by Tivoli Storage Productivity Center. However, when you click the magnifying glass icon next to a row in that report, the resulting report displays disk totals for

computers only. This action might cause the totals displayed in the network-wide report to appear larger than the totals reflected in the report that you drill down in to (if you have disks that are part of both computers and storage subsystems within your network).

#### **Report values for General Parallel File Systems (GPFS) file systems**

For reports that display information collected by scans, Tivoli Storage Productivity Center does not include statistics for GPFS snapshot directories and the files in those directories. Because of this issue, the values in these columns might appear smaller than the actual numbers for the file systems.

For example, when viewing data collected by scans of GPFS file systems in the Data Manager > TPC-wide Storage Space > File System Space report, the values for the File Count, Directory Count, and File System Space columns might be lower than the actual numbers for the file systems.

#### **IBM XIV Storage System reporting limitation for disk drives**

Tivoli Storage Productivity Center does not collect physical serial number information about the disk drives that are associated with the XIV system devices. Therefore, any reports that display information about the XIV system do not include the physical serial numbers of the disk drives.

#### **Hitachi Data Systems storage subsystems limitations**

- There is a limitation within reports when Tivoli Storage Productivity Center collects information about Hitachi Data Systems (HDS) storage subsystems that are configured as storage virtualizers. Specifically, any reports that display a value for the Disk Unallocated Space column do not count the storage consumed by volumes residing on external disks when the HDS is used as a virtualizer.
- Tivoli Storage Productivity Center does not support Logical Unit Size Expansion (LUSE) with Hitachi storage subsystems. Because of this limitation, Tivoli Storage Productivity Center might report a different number of volumes for a subsystem than is reflected in the HiCommand Device Manager (HDvM) when LUSE is used.

#### **Time zones for computers monitored by Data agents and Storage Resource agents**

The time zones of computers that are monitored by Storage Resource agents are shown as Greenwich Mean Time (GMT) offsets in Tivoli Storage Productivity Center reports. The time zones of computers that are monitored by Data agents are shown as the local time zone in Tivoli Storage Productivity Center reports. For example, a computer located in Los Angeles shows the following time zones in the By Computer report in Asset reporting:

- If monitored by a Storage Resource agent: (GMT-8:00) GMT-8:00
- If monitored by a Data agent: (GMT-8:00) America/Los\_Angeles Time zone

#### **SQL access to data collected by Tivoli Storage Productivity Center for use in external reports and applications**

You can use the data collected by data collection jobs to create custom reports in external tools. Tivoli Storage Productivity Center provides a set of DB2<sup>®</sup> views that represent key data that has been stored in the database repository. Use Structured Query Language (SQL) commands to retrieve the data in these views. Use this data in external tools such as Business Intelligence and Reporting Tools (BIRT), Microsoft Excel, and Tivoli Common Reporting. See *Planning to retrieve information from views in the database repository* topic in the *IBM Tivoli Storage Productivity Center and IBM*

*Tivoli Storage Productivity Center for Replication Installation and Configuration Guide* for more information about how to use SQL to retrieve data for use in custom reports and tools.

### **Storage Pools in EMC storage subsystems**

The data that Tivoli Storage Productivity Center collects about EMC storage subsystems is based on SMI-S standards. Within SMI-S model, different types of storage pools are defined and several subsystem attributes can be modeled as special pool types. Beside the pools that are available for volume creation, additional storage space from the EMC storage subsystem might be represented as a pool within the SMI-S model. Because of this representation, Tivoli Storage Productivity Center reports and CLI commands might display more pools for EMC storage subsystems than what is reported through the EMC Navisphere agent.

### **Computers in a Windows domain**

Once computers are discovered in a domain, they remain in the database until three discoveries do not report them. If you remove the computer that is doing the discovery in a domain, no further discovery results are processed by the server. Therefore, any remaining domain computers that were previously discovered will continue to show their unreachable value in Unmanaged Computers reports.

### **UDB reports**

For UDB reports, the "User Name" and "Creator" columns represent the schema (not the operating system user) which is using the storage space. Each schema is associated with one operating system user. The reports show usage at the schema level.

### **Reports display incorrect number of folders for Windows computers**

On monitored computers that run Windows, the number of folders that is reported by Tivoli Storage Productivity Center might be different from the number of folders that is indicated by Windows Explorer. This situation occurs when settings on the computer indicate that files, folders, and drives are hidden. For example, the Data Manager > Reporting > Usage > files > File Summary > by Computer report might display a number of folders for a Windows computer that is less than the number of folders provided by Windows Explorer on that computer.

To ensure that the number of folders reported by Tivoli Storage Productivity Center more closely matches the number reported by Windows Explorer, complete the following steps:

1. On each of the Windows computers that you are monitoring, start Windows Explorer.
2. Navigate to the file system whose settings that you want to change.
3. Select **Folder and search options...** or **Folder options...** from the **Organize** or **View** menus.
4. Click the **View** tab.
5. In the **Advanced Settings** section, make the following changes:
  - Select **Show hidden files, folders, and drives**.
  - Clear **Hide protected operating system files**.
6. Run probe and scan data collection schedules on the computers.
7. View the reports that display the number of folders for Windows computers.

### **Information about space efficient volumes**

Keep in mind the following information about space-efficient volumes:

- A volume name or virtual disk in a report that is marked with an asterisk (\*) indicates that the volume or virtual disk is a space-efficient volume. The \* is not part of the actual volume or virtual disk name.
- Tivoli Storage Productivity Center does not include space efficient volumes when calculating the values for consumed and available space on DS8000 storage systems. Because of this, the values for these columns in Tivoli Storage Productivity Center reports might be different than the space that is actually available on DS8000 storage systems with space efficient volumes.
- Tivoli Storage Productivity Center does not display information about space-efficient volumes for IBM DS8000 storage systems in the following reports:
  - Disk Manager > Reporting > Storage Subsystems > Computer Views > By Filesystem/Logical Volume
  - Disk Manager > Reporting > Storage Subsystems > Storage Subsystem Views > Disks

**The mapping of volumes to an ESX hypervisor might not display correctly**

When assigning and unassigning volumes to an ESX hypervisor, the mapping of the volumes might not display correctly in Tivoli Storage Productivity Center reports. For example, if you assign a volume to a hypervisor, unassign it, and then assign a different volume to the hypervisor, information about the original volume might still be displayed in the Correlated Storage Subsystem volume report that is available under Data Manager > Reporting > Asset > By Hypervisor.

To ensure that Tivoli Storage Productivity Center reports display the correct volume information, complete the following tasks before generating the reports:

- After unassigning a storage system volume from a hypervisor, use VMware vSphere Client to rescan the HBA.
- After assigning a new storage system volume to the hypervisor, use VMware vSphere Client to rescan the HBA.
- In Tivoli Storage Productivity Center, rerun probes for the computer, hypervisor, and storage system.

---

## Working with reports

Learn how to drill down through reports, navigate report tabs, view historical trends in reports, and view reports as tables of data or as graphical charts. The ability to customize your view of reports provides you with the flexibility to get a very detailed look at your information or a very high level overview.

## Generating and viewing reports

Learn how to generate and view reports.

Each of the report types available in IBM Tivoli Storage Productivity Center are generated in different ways to provide you with the flexibility to meet your environment's changing needs. Click the following links to learn more about how to generate the different report types and determine what works best for your organization.

| Report type                 | How to generate                  |
|-----------------------------|----------------------------------|
| User-defined online reports | "Generating reports" on page 793 |

| Report type             | How to generate                                     |
|-------------------------|-----------------------------------------------------|
| System reports          | "Generating and viewing system reports" on page 980 |
| User ID (saved) reports | "Creating a user-defined report" on page 1035       |
| Batch reports           | "Creating batch reports" on page 1021               |
| Rollup reports          | "Generating rollup reports" on page 1036            |

## Sorting report columns

Learn how to sort the columns in reports.

Reports appear in a tabular format and are comprised of rows and columns. You can scroll the report up and down and to the right and left to view the entire report. Use the **View** menu to hide/show the navigation tree to increase the viewable area of the report. You can also drag the divider bar in the middle of the screen back and forth to reallocate the amount of space that is given each pane.

**To sort the columns or change the order of the columns in a report:**

- **Re-sorting columns.** To resort a column, click the column header. The column will resort in ascending or descending order.
- **Moving columns.** To change a column order, press <Alt>, and then click and drag the column you want to move.

## Printing reports or charts


Learn how to print reports to a printer, a PDF, an HTML file, a CSV file, or a formatted text file.

You can print any IBM Tivoli Storage Productivity Center report or chart displayed in the content pane.

To print a report or chart, complete the following steps:

1. Generate the report you want.



2. Click **File > Print** or click  on the tool bar. A print window opens.

3. Select the output that you want:



- Select **Printer** to send the output directly to a printer
- Select **PDF file** to save the output as a PDF file
- Select **HTML file** to save the output as an HTML file
- Select **CSV file** to save the output as a comma delimited file. Select **Include Headers** to include column headers in the report output file. Select **Include Totals** to include the totals column in the report output file.
- Select **Formatted file** to save the output as a formatted text file

4. Click **Continue**.

- If you selected **Printer**, the Page Setup window opens. Specify any custom settings for the print job on the Page Setup window and click **OK**. The report or chart is printed to your default printer.
- If you selected **PDF file**, **HTML file**, **CSV file** or **Formatted file**, the Save As window opens. Select the location where you want to store the file and enter a name for the file in the **File Name** field. Click **Save**. Depending on the file type you selected, the report or chart is saved as a PDF or HTML file.

## Including charts in reports

Learn how to create charts from the reports that you generate.

There are a large number of charts and chart types that can be generated from reports. All reports that support charting will have a  button at the top of the report. Any row in a report that supports charting on that row will have a  beside it.

Many charts can be generated simultaneously from the same report. After generating a chart, return to the **Report** tab and choose another chart or another set of objects to generate a chart upon. Another **Chart** tab will be generated to the right of the already existing tabs and you will be positioned to the new chart tab.


See “Generating charts” for information on how to generate charts from reports.

You can also customize the appearance of charts to determine how you want to present data. For example, you can create pie chart, bar chart, stacked bar chart, determine how many charts appear per page, specify time increments, and so on. See “Customizing charts” on page 787 for information on how to customize charts.

### Generating charts

Learn how to generate charts.


#### To view a chart of all of the rows in the table:

1. Click , which is at the top of the report. All reports that support charting will have this multi-colored pie chart icon at the top of the report.
2. Select the type of chart you would like to generate. The types of charts available will vary depending on the report. After you select the chart type, a new tab for that chart will appear in the content pane.


#### To generate a chart for just the element in that row:

- Click the chart icon that appears to the left of that report row.

#### To generate a chart for multiple, non-consecutive rows in a report:

1. Select rows in the report by holding down the control key (Ctrl key) on the keyboard while clicking the desired rows.
2. Click any  button that appears to the left of the selected rows and select the type of chart you would like to generate. After you select the chart type, a new tab for that chart will appear in the content pane.

#### To generate a chart for a range of rows in a report:

1. Click the first row in the range and then hold down the Shift and Ctrl while clicking the end row in the range.
2. Click any  button that appears to the left of the selected rows and select the type of chart you would like to generate. After you select the chart type, a new tab for that chart will appear in the content pane.






Figure 53. Chart Generation

## Customizing charts

Learn how to customize the appearance of charts.

While viewing a chart, you can customize its appearance using the Chart Customization window. This includes:

- Entering the title that appears at the top of a chart
- Determining the maximum number of charts or series that appear per screen
- Specifying how to show values for data points or areas
- Selecting a type for a chart such as a pie, bar, stacked bar, and 100% stacked bar
- Determining the orientation (vertical or horizontal) of a chart
- Specifying history chart settings such as trends and time increment values

To access this window from a chart, right-click the chart and select **Customize this chart** from the pop-up menu. The Chart Customization window appears.

Use this window to customize the chart generated from a report, a report created in the Reporting facility, a batch report, user reports, and dashboard reports. You can use the settings on this window to:

- Customize the display of the current chart.
- Customize the display of the current chart and save those changes. If you save the customization settings for a report chart, that chart will use those settings whenever it is generated.

If a field on this window does not apply to a selected report it will appear as grayed out and cannot be edited.



| Field                                         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Chart title                                   | Enter the name you want to appear at the top of the generated chart.                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Maximum number of charts or series per screen | <p>Select the maximum number of charts or series per page that should appear in the generated report from this list. For example, how many pie charts will appear on one page. When the number of charts that appear on one page exceeds the number in this field, the additional charts will be shown on a separate page.</p> <p>The most charts or series you can define to appear on a page is 30.</p>                                                                                       |
| Show values for data points or areas          | <p>Determine how to display the data points/areas on a chart. You can choose to display the data points areas:</p> <ul style="list-style-type: none"> <li>• Only when hovering - display values for data points on an online chart when the mouse hovers over an appropriate place on that chart</li> <li>• Always - always display and print values for data points</li> </ul>                                                                                                                 |
| Graph charts                                  | Use the fields in this section to customize the appearance of graphs and charts.                                                                                                                                                                                                                                                                                                                                                                                                                |
| Chart type                                    | <p>Select a type for the generated chart including:</p> <ul style="list-style-type: none"> <li>• Pie Chart</li> <li>• Bar Chart</li> <li>• Stacked Bar</li> <li>• 100% Stacked Bar</li> </ul>                                                                                                                                                                                                                                                                                                   |
| Chart orientation                             | Determine whether the chart graphic will appear vertically or horizontally on a page. You can set the orientation for the following chart types only: Bar Chart, Stacked Bar, and 100% Stacked Bar.                                                                                                                                                                                                                                                                                             |
| Maximum number of series per graph            | <p>Select the maximum number of series that should appear in one graph or chart. For example, how many slices will appear in one pie chart. When the number of series that appear in one graph exceeds the number in this field, the additional charts will be shown on a separate page.</p> <p>The default maximum is 12. The most series that can appear in one chart is 30.</p> <p><b>Note:</b> Disk Capacity batch reports always display one series per chart, and one chart per page.</p> |
| History Charts                                | Use the fields in this section to customize the appearance of history charts.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Show Trends                                   | Check this box if you want IBM Tivoli Storage Productivity Center to show a trend for the data in a history chart. The trend attempts to show future storage usage in a chart based on earlier usage.                                                                                                                                                                                                                                                                                           |

| Field                     | Description                                                                                                                                                                                                                               |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Time increment for x-axis | Select which value to use for the time increment that appears on the x-axis of a history chart. You can select By Scan, Weekly, Monthly. For example, if you choose Monthly, the x-axis will measure storage usage in monthly increments. |
| Restore Defaults          | Reset the values in the fields on this window to their original default settings.                                                                                                                                                         |
| Apply                     | Apply the settings defined in this window to the current chart.                                                                                                                                                                           |
| Save                      | Apply the settings defined in this window to the current chart and save those settings.<br><br>When the chart is generated in the future, it will automatically use the saved settings.                                                   |
| Cancel                    | Exit the window without applying or saving the settings in the window.                                                                                                                                                                    |

## Drilling down a report row

Learn how to drill down into report rows to view more detailed information about your storage resources.

Many of the rows that are presented in the generated reports allow you to drill down and get more reports of more detailed information. For example, if you are looking at a report on computers, you might be able to drill down on a row for a specific computer and generate a report for the file systems contained on that computer.

When detailed drill down is allowed within a report, there is a  icon at the left of that row. To drill down, click  and a new tab with the drill down report will be generated and added to the tab dialog. You will then be positioned on that tab viewing the new detail report that you have generated.

## Navigating tabs

Learn how to navigate report tabs to view detailed sub reports.

As you generate reports and drill down on rows to generate detailed sub reports, the number of tabs in the Reporting window will grow. You can move about from tab to tab by clicking the tab name.

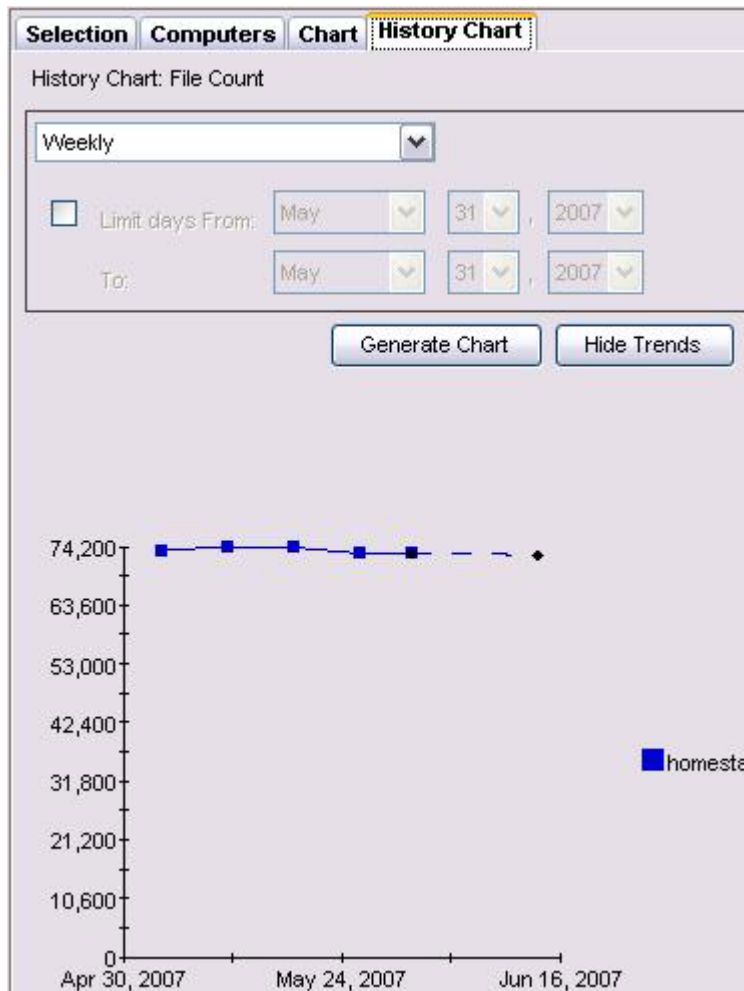


Figure 54. Tabs

You can also return to the **Selection** tab and change your selection criteria to start all over again with a newly generated report at any time. If you change your selection criteria and generate a new report, all subsequent or dependent tabs (for example, drill down and chart tabs) are closed.

To close a tab, right-click the tab folder top and select **Close** from the pop-up menu.

## Trending with historical data

Learn how to view trends related to the capacity and usage of your storage resources.

Use the trending of historical data to view patterns of your data across your entire network. For example, if you set up a graph that shows the disk capacity of a computer between January 1 and December 31, you can view "trended" data that will predict the disk capacity for that computer over the next 6 months. The trended data appears as a dashed line that is appended to the last probed or scanned data point.

The length of time for how far in the future that trended data will show is determined by 50% of the length of time displayed in the graph. For example, if

the data in the non-trended portion of the graph displays data between January 1 and December 31, the predicted trend for that data will be shown for the next 6 months (January 1 to December 31 = 12 months, 50% of 12 months = 6 months trended). You can change the amount of time trended by changing the length of time of the graphed data using the **Limit days from** and **To** fields.

Before viewing trends in the data collected by IBM Tivoli Storage Productivity Center, use the history aggregator function to define and run jobs to sum data in the database repository for historical reporting purposes. For example, after running the history aggregator, you can view the sum of usage across multiple resources, by database.

To view a historical data, generate a report, right-click the row or rows for which you want to view the data, and select an option from the menu.

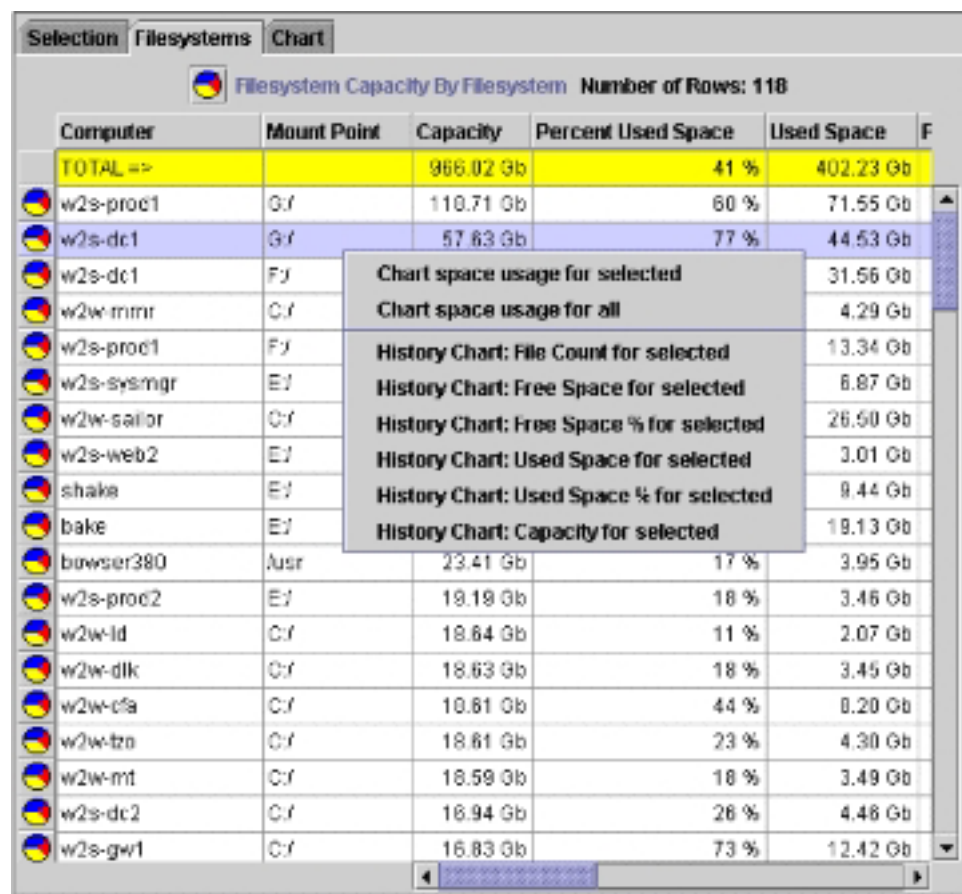


Figure 55. Trending with Historical Data

Depending on the type of report you generated, you can create the following history charts for the storage resources that you select:

| History chart             | Description                                                                                |
|---------------------------|--------------------------------------------------------------------------------------------|
| Free Space for Selected   | Shows the free space available over a specified period of time. <sup>1</sup>               |
| Free Space % for Selected | Shows the percentage of free space available over a specified period of time. <sup>1</sup> |
| Used Space for Selected   | Shows used space over a specified period of time. <sup>1</sup>                             |

| History chart                  | Description                                                                                                                                                                                                                               |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Used Space % for Selected      | Shows percentage of used space over a specified period of time. <sup>1</sup>                                                                                                                                                              |
| Capacity for Selected          | Shows capacity for selected over a specified period of time. <sup>1</sup>                                                                                                                                                                 |
| Unallocated Space for Selected | Shows the unallocated space available for the selected resources over a specified period of time. <sup>1</sup>                                                                                                                            |
| Unallocated Space for All      | Shows the unallocated space available for all the storage resources in the report over a specified period of time. <sup>1</sup>                                                                                                           |
| File Count                     | Shows the number of files accessed or modified over a period of: <ul style="list-style-type: none"> <li>• 1 day</li> <li>• 1 day - 1 week</li> <li>• 1 week - 1 month</li> <li>• 1 month - 1 year</li> <li>• &gt; year</li> </ul>         |
| File Count %                   | Shows the percentage of files accessed or modified over a period of: <ul style="list-style-type: none"> <li>• 1 day</li> <li>• 1 day - 1 week</li> <li>• 1 week - 1 month</li> <li>• 1 month - 1 year</li> <li>• &gt; year</li> </ul>     |
| Size                           | Shows the size of files accessed or modified over a period of: <ul style="list-style-type: none"> <li>• 1 day</li> <li>• 1 day - 1 week</li> <li>• 1 week - 1 month</li> <li>• 1 month - 1 year</li> <li>• &gt; year</li> </ul>           |
| Size %                         | Shows the percentage of file size accessed or modified over a period of: <ul style="list-style-type: none"> <li>• 1 day</li> <li>• 1 day - 1 week</li> <li>• 1 week - 1 month</li> <li>• 1 month - 1 year</li> <li>• &gt; year</li> </ul> |

<sup>1</sup> The time options for data points in a chart are **By Aggregator**, **Weekly**, and **Monthly**. You can select one of these options from the list menu on the **History Chart** page for the report. (See Figure 56 on page 793.)

- Select **By Aggregator** to generate a chart in which its data points correspond to each time a history aggregator job is run. By default, the history aggregator runs daily at 3:30 a.m., so when you generate a history chart with this default setting there is a point in the chart for every day in the selected time range. If you change the schedule of the history aggregator, the points in the history chart that you generate change accordingly.
- Select **Weekly** to generate a chart that includes one data point for each week in the selected time range.
- Select **Monthly** to generate a chart that includes one data point for each month in the selected time range.

Figure 56 on page 793 shows a history chart for used space for a file system:

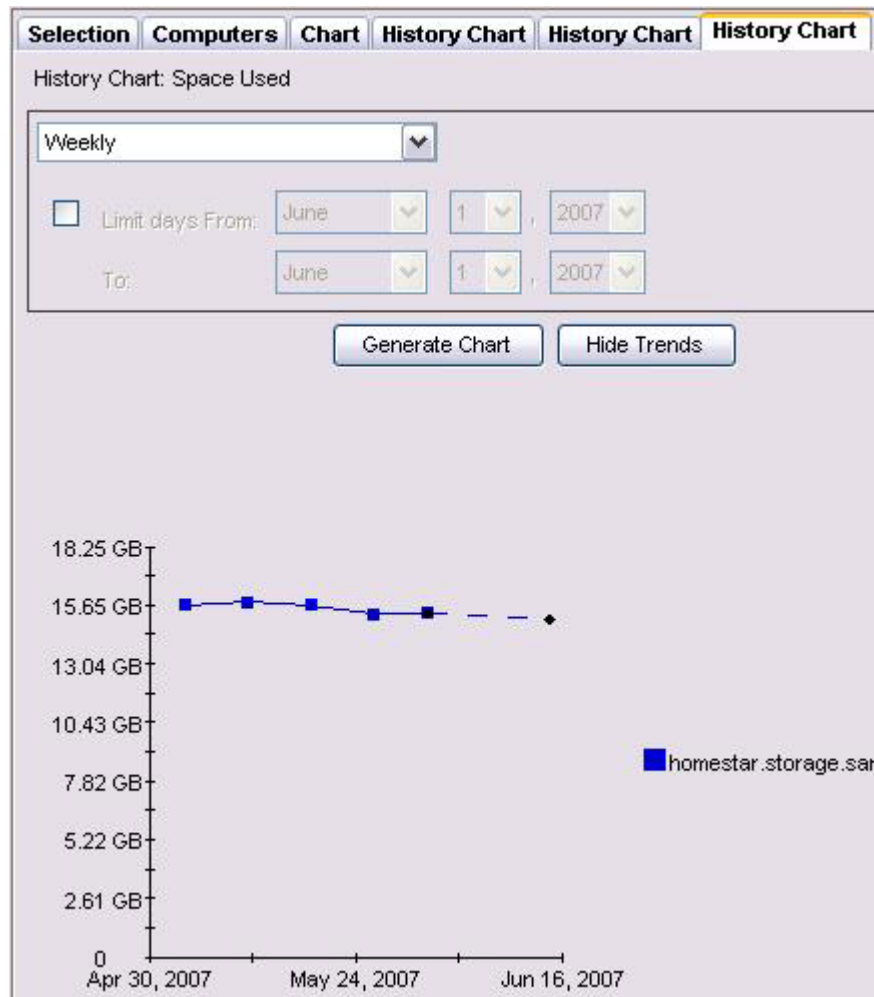


Figure 56. History Chart for File System Used Space

Select **File** → **Export data** to export history data to a comma delimited file (.csv). You can include or not include headers in the comma delimited file that you create. The column headings within the header reflect the type of report from which you are exporting history data.

## User-defined online reports

Use the information in this section to learn how to generate and view online reports. These reports are available in the **Reporting** nodes under Data Manager, Data Manager for Databases, Disk Manager, and Fabric Manager.

### Generating reports

Learn how to generate and view online reports in Data Manager, Data Manager for Databases, Disk Manager, and Fabric Manager.

1. Expand the navigation tree to display the type of report you want to generate. For example, if you want to generate a disk capacity report, expand **Data Manager > Reporting > Capacity > Disk Capacity**.
2. Highlight a By Clause to determine how you want to view a report. For example, to view capacity information according to computer group, click **Reporting > Capacity > Disk Capacity > By Computer Group**. The **Selection** page is displayed.

3. Use the **Selection** page to select the profile to use when generating a report and determine what columns appear within a report. See “Selecting resources to appear on a report” on page 796 for information.
4. Click **Selection...** to select the objects that you want to report upon from the Select Resources window.
5. Click **Filter...** to further filter the objects that appear in a report. Filters enable you to apply general rules to the report based on the rows in that report. See “Filtering the resources that appear on a report” on page 799 for more information on filtering.
6. Use the **Available Columns** and **Included Columns** list boxes to determine what columns are displayed in a generated report.
7. Click **Generate Report**. A new tab will be added to the tab dialog representing the report that you generated.
8. Click the new tab to view the report. Reports are tabular in format and composed of rows and columns. You can scroll the report up and down and to the right and left to view the entire report. Use the **View** menu to hide/show the navigation tree to increase the viewable area of the report or drag the divider bar in the middle of the screen back and forth to reallocate the amount of space that is given each pane.

### Defining options for a report

Learn how to customize a report before generating it.

Depending on the report that you want to generate and the level upon which you have chosen to report (for example, by computer, by filesystem, network wide, and so on), you are typically presented with a **Selection** page. For example, if you want to generate the Most Obsolete Files report, expand **Data Manager > Reporting > Usage > Files > Most Obsolete Files > By Computer**. Use the **Selection** page to:

- Select the profile to use when generating a report
- Select the resources (for example, computers, computer groups, storage subsystems, and so on) to display in a generated report
- Apply filters to the columns within a report (for example, show only the systems that contain over 80% used space) to determine what resources appear within the rows of a report
- Determine what columns appear within a report

The following table describes the interface elements that might appear on the **Selection** page (depending on the report that you want to generate):

| Interface element                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Profile</b><br>(Data Manager only) | This field indicates the profile to use when generating the report. If you just want to use the supplied default profile, then leave the field as is and generate the report. However, if you want to use a customized profile that you have set up for this reporting, choose that profile from the list.<br><br>Keep in mind that profiles control what statistics are gathered and also what files are included in that statistical analysis. |



| Interface element                                                                                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Return maximum of __ rows per <i>storage_resource</i>                                            | <p>This field controls how many rows appear in a generated report. The value that appears as the <i>storage_resource</i> in this field depends on the selected By clause. For example, if you select <b>Data Manager &gt; Reporting &gt; Usage &gt; Files &gt; Most Obsolete Files &gt; By Computer</b>, you can enter <b>10</b> in the <b>Return maximum of __ rows per computer</b> field and to limit the number of rows in the report to only 10 obsolete files per computer. If you have set up scans to collect many more rows than that per computer, this is an effective and efficient way to control the size of the generated report.</p> <p>This field enables you to control the size of the report that is generated for different by clauses.</p> |
| Limit to this time range:                                                                        | The fields in this section determine the date range from which to gather availability information for the objects that you select.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Relate Computers to:</b><br>(Disk Manager - Storage Subsystem reports only)                   | <p>This field generates a report that shows the relationship between the computer and one of these options: storage subsystems, volumes, or disks.</p> <ul style="list-style-type: none"> <li>• Storage Subsystem. Choose this option to view the Storage Subsystems that the logical volumes within a computer are using.</li> <li>• volumes. Choose this option to show the volumes within a storage subsystem that the logical volumes within a computer are using.</li> <li>• Disks. Choose this option to show the disks within a disk array that the logical volumes within a computer are using.</li> </ul>                                                                                                                                               |
| <b>Relate Filesystems/Logical Volumes to:</b><br>(Disk Manager - Storage Subsystem reports only) | <p>This field generates a report that shows the relationship between the file systems/logical volumes and one of these options: storage subsystems, volumes, or disks.</p> <ul style="list-style-type: none"> <li>• Storage Subsystem. Choose this option to view the storage subsystems on which a file system/logical volume is allocated.</li> <li>• volumes. Choose this option to view the volumes on which the file system/logical volume is allocated.</li> <li>• Disks. Choose this option to show the disks on which a file system/logical volume is allocated.</li> </ul>                                                                                                                                                                              |
| <b>Relate Storage Subsystems to:</b><br>(Disk Manager - Storage Subsystem reports only)          | <p>This field generates a report that shows the relationship between a storage subsystem and one of these options: computer or file system/logical volume.</p> <ul style="list-style-type: none"> <li>• Computer. Choose this option to view the computers that are using a storage subsystem.</li> <li>• Filesystem/Logical Volume. Choose this option to view the file systems/logical volumes that are using a storage subsystem.</li> </ul>                                                                                                                                                                                                                                                                                                                  |
| <b>Relate Volumess to:</b><br>(Disk Manager - Storage Subsystem reports only)                    | <p>This field generates a report that shows the relationship between a volume and one of these options: computer or file system/logical volume.</p> <ul style="list-style-type: none"> <li>• Computer. Choose this option to view the computers that are using a volume.</li> <li>• Filesystem/Logical Volume. Choose this option to view the file systems/logical volumes that are using a volume.</li> </ul>                                                                                                                                                                                                                                                                                                                                                   |
| <b>Relate Disks to:</b><br>(Disk Manager - Storage Subsystem reports only)                       | <p>This field generates a report that shows the relationship between a disk and one of these options: computer and file system/logical volume.</p> <ul style="list-style-type: none"> <li>• Computer. Choose this option to view the computers that are using a disk.</li> <li>• Filesystem/Logical Volume. Choose this option to view the file systems/logical volumes that are using a disk.</li> </ul>                                                                                                                                                                                                                                                                                                                                                        |

| Interface element                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Selection...                         | <p>This button is used to determine the storage resources (computers, file systems, logical volumes, storage subsystems, and so on) upon which you want to report. When you click <b>Selection...</b>, the Select Resources window opens. If you clear the check box for a resource on the Select Resource window, a push-pin icon displays on the <b>Selection...</b> button.</p> <p>If you are generating a Network-wide report, the <b>Selection...</b> button will be disabled because you have chosen to generate the report network wide, for all objects.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Filter...                            | <p>Use this button to further select or restrict the values within the columns you want to include on a report. Filters enable you to apply general rules to the selected columns. When you click <b>Filter...</b>, the Edit Filter window opens.</p> <p>For example, select all computers with the name beginning with <i>w2s-*</i>, or select all file systems with percent used space <math>\geq</math> to a specific value.</p> <p>If a filter is defined for this report, a pushpin icon displays on the <b>Filter...</b> button.</p> <p><b>Note:</b> For storage entities that contain an asterisk (*) within their names, the * is used as a wildcard character if you select <b>LIKE</b> from the <b>Operator</b> menu on the Edit Filter panel. See “Filtering the resources that appear on a report” on page 799 for more information.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Available Columns / Included Columns | <p>These panels display the columns that can be included or are already included in the generated report.</p> <ul style="list-style-type: none"> <li>• The <b>Available Columns</b> list displays the columns that you can include in the generated report that are not already included.</li> <li>• The <b>Included Columns</b> list displays the columns that will appear in the generated report. By default, all the columns for a report are listed in this list box and will be displayed in the generated report. You can highlight the columns and use the up and down arrows to rearrange columns in the report. The order of the items in this list determines the order in which the columns will appear in the generated report. The greyed-out names shown at the top of the list are frozen columns of the report. They cannot be selected, removed, or reordered.</li> </ul> <p>To include a column in a report:</p> <ol style="list-style-type: none"> <li>1. In the <b>Available Columns</b> list, click the name of the column. To select additional columns, press Shift and click the columns that you want to add or press Ctrl and click the first and last column to select a range of columns.</li> <li>2. Click the right-arrow button (&gt;&gt;) to add the columns to the <b>Included Columns</b> list.</li> </ol> <p>To remove a column in a report:</p> <ol style="list-style-type: none"> <li>1. In the <b>Included Columns</b> list, click the name of the column. To select additional columns, press Shift and click the columns that you want to add or press Ctrl and click the first and last column to select a range of columns.</li> <li>2. Click the left-arrow button (&lt;&lt;) to add the columns to the <b>Available Columns</b> list.</li> </ol> |

## Selecting resources to appear on a report

Learn how to select the storage resources that will appear in a report.

Use the Select Resource window to determine the objects (computers, filesystems, logical volumes, storage subsystems, etc.) upon which you want to report. To select the resources on which you want to report:

1. Select the type of report you want to generate. For example, **Data Manager > Reporting > Capacity > Disk Capacity > By Disk**.
2. Click **Selection**. The Select Resources window opens.

The objects that appear in this window are determined by the report you are generating and the By Clause you selected. For example, if you choose to generate the report **Data Manager > Reporting > Capacity > Disk Capacity > By Disk**, the Select Resource window will display a list of computers and subsystems that you can include in the report.

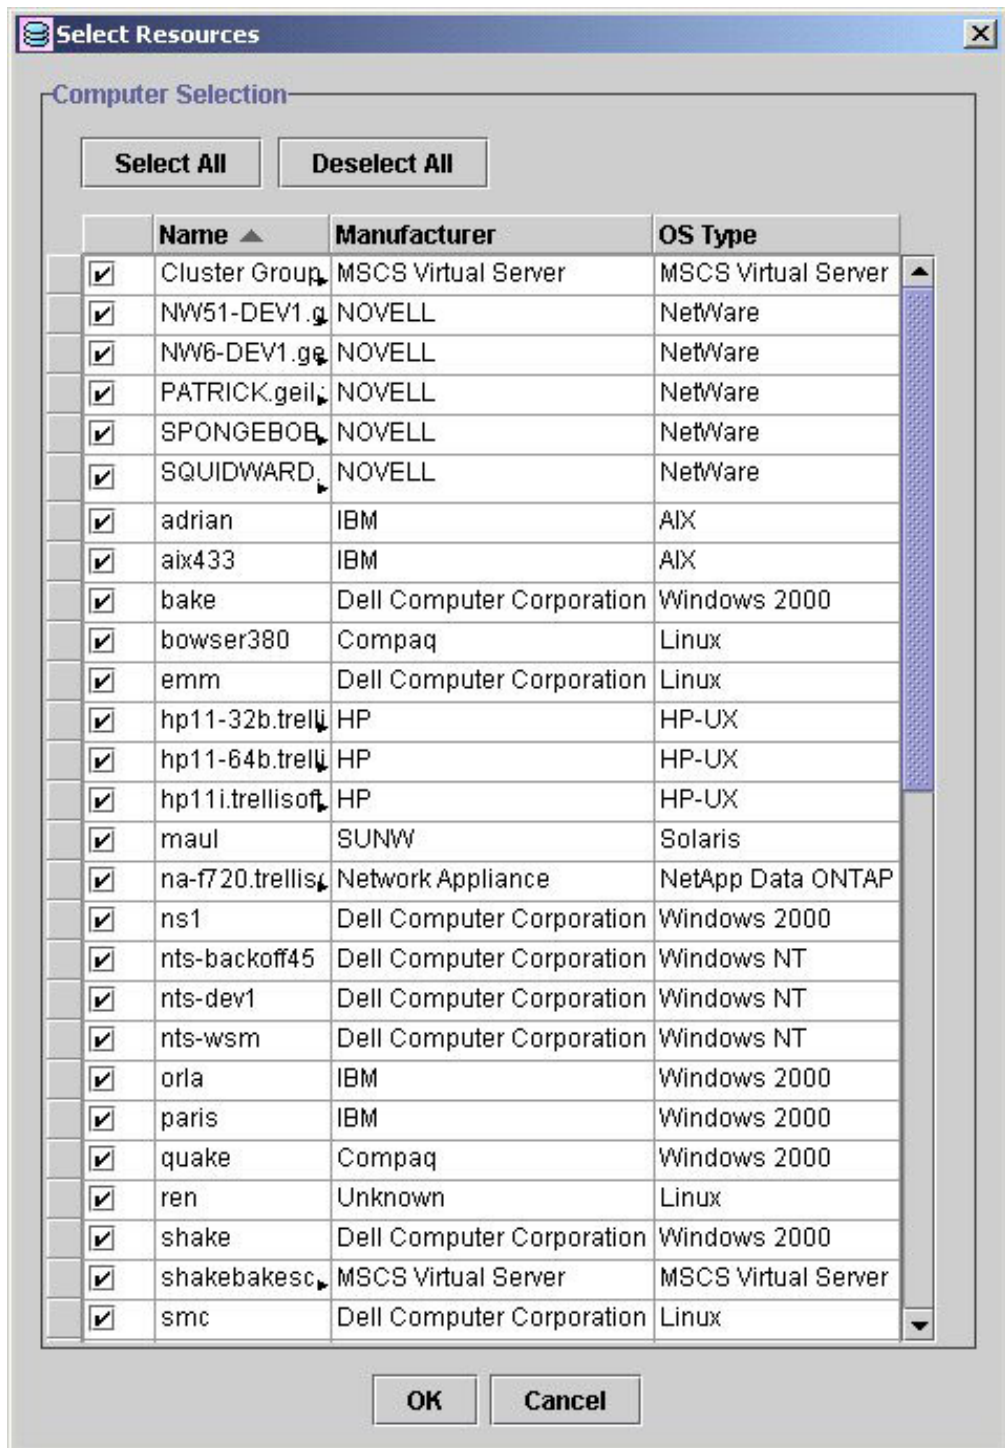


Figure 57. Selection Resources window

When the Select Resource window is initially opened, all objects in the list are selected.

- If you would like to generate the report for all the objects that have been selected, click **OK**.
- To select/deselect individual objects in the selection list, simply select the check box to the left of the object. When there is a check mark in the check box, then the object has been selected for inclusion in the report.

You can use **Select All** and **Deselect All** to quickly select and deselect all of the objects in the list.

**Note:** If you are generating a Network-wide report, the **Selection** button is disabled on the **Selection** tab because you have chosen to generate the report network wide, for all objects.

## Filtering the resources that appear on a report

Learn how to apply conditions to the columns in a report to further filter the resources that appear in a report.


Use the Edit Filter window to apply conditions to the columns within a report to further select or restrict objects on which you want to report. Filters enable you to apply general rules to the report. For example, you can show only the file systems that contain over 80% used space, or show computers that match "w2s-\*", or select all disks with percent used space  $\geq$  to a specific value. To filter the objects that appear in a report:

1. Select the type of report you want to generate. For example, **Data Manager > Reporting > TPC-wide Storage Space > Disk Space > By Disk**.
2. Click **Filter....** The Edit Filter window opens.

**Note:** For storage entities that contain an asterisk (\*) within their names, the \* is used as a wildcard character if you select **LIKE** from the **Operator** menu on this panel. For example, if you enter \*XYZ in the **Value 1** field and select **LIKE** from the **Operator** menu, the resulting report includes any entities named "\*XYZ", and also all entities whose names end with the letters "XYZ". To create a filter that includes only the storage entity named "\*XYZ" in the resulting report, enter \*XYZ in the **Value 1** field and select **IN** from the **Operator** menu.

The following table describes the interface elements that appear on the **Edit Filter** window:

| Interface element                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Add</b>                               | This button is used to specify multiple conditions for the filter. When you click <b>Add</b> , report columns will be displayed enabling you to define conditions for those columns.                                                                                                                                                                                                                                                        |
| <b>Delete All</b>                        | This button is used to delete all the conditions displayed on the table.                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Case-sensitive string comparisons</b> | This check box indicates that case-sensitive string comparisons should be done. Select this check box to ensure that values appearing in the generated report match the case of the values entered in the Value 1 and Value 2 fields.                                                                                                                                                                                                       |
| <b>Records must meet</b>                 | This field specifies how the conditions in the filter are grouped together. <ul style="list-style-type: none"><li>• <b>All conditions.</b> All of the entered conditions (rows on the dialog box), must be true in order for the record to show up as a row in the report.</li><li>• <b>At least one condition.</b> At least one of the entered condition must be true in order for the record to show up as a row in the report.</li></ul> |

| Interface element       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Column</b>           | <p>This field is used to select specific columns on which you want to define conditions. The list on this field shows the columns available for the report. As you select a column from the list, it is displayed on the table, and the column name is removed from all other lists.</p> <p>Click</p>  <p>to delete the column.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Operator</b>         | <p>This field is used to select the condition's operator. Not all operators apply to all the column data-types.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Value 1/ Value 2</b> | <p>This field is used to specify fixed value(s) to be compared against the record's column value.</p> <p>Depending on the operator and data type, this field may be absent or display-only. If the field is display-only, an <b>Edit...</b> button will be enabled. If both this field and the <b>Edit...</b> button are enabled, this field allows shorthand value entry. If you click <b>Edit</b> , a window opens to provide more assistance in editing the value.</p> <p>A <b>Value 2</b> field will be present only for (NOT) BETWEEN.</p> <p>For certain columns that contain string text, you can:</p> <ul style="list-style-type: none"> <li>• put the string in quotes for the <b>LIKE</b> operator, or</li> <li>• use the = operator without quotes instead of <b>LIKE</b> (the default)</li> </ul> <p>For example, if you want to filter the Device Path (computer) column on the <b>Reporting → Storage Subsystems → Computer Views → By Filesystem/Logical Volume</b> report, you can search/filter on a name such as "Disk 1" or "Disk 2" with the <b>LIKE</b> operator. For the = operator, you can enter <b>Disk 1</b> or <b>Disk 2</b> with no quotes. <b>Note:</b> An asterisk is prepended to the names of space-efficient VDIsks in the IBM Tivoli Storage Productivity Center database repository. When you use a filter to determine which volume names appear in a report, you must include the "*" as the first character in the filter value to ensure that those VDIsks are included. For example, if you want to include VDIsks whose name begins with "vdisk1", you must enter <b>*vdisk1</b> in the <b>Value 1</b> field and select <b>LIKE</b> or <b>IN</b> from the <b>Operator</b> menu. If you select the <b>LIKE</b> operator, the *vdisk1 VDisk is included in the report, as well as other VDIsks whose names end in "vdisk1".</p> |
| <b>Edit</b>             | <p>This button is used to edit values pertaining to size, date, and for specific values used for IN, LIKE, etc. operators. Depending on the condition you are creating, a dialog box displays to help you to enter values that correspond to the condition.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

| Interface element | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Selection...      | <p>This button is used to determine the objects (computers, file systems, logical volumes, storage subsystems, etc.) upon which you want to report. The Select Resources window opens. If you clear the check box for a resource on the Select Resource window, a pushpin icon displays on the <b>Selection...</b> button.</p> <p>If you are generating a Network-wide report, the <b>Selection...</b> button will be disabled because you have chosen to generate the report network wide, for all objects.</p> |
| OK                | <p>This button is used to save the filter. You are returned to the <b>Selection</b> tab. The <b>Filter</b> button will display a pushpin to indicate that a filter has been defined for the report.</p>                                                                                                                                                                                                                                                                                                          |
| Cancel            | <p>This button is used to return to the <b>Selection</b> tab without saving the filter.</p>                                                                                                                                                                                                                                                                                                                                                                                                                      |

## Using reporting groups

A reporting group is a user-defined set of storage assets, monitoring groups, and reporting groups upon which you want to report.

Use reporting groups to view storage information about:

- Resources that appear in different monitoring groups. Because a resource can appear in only one monitoring group at a time, you can use reporting groups to include resources from different monitoring groups and view reports on them.
- A subset of resources within a monitoring group. For example, you can create a reporting group for a specific subset of computers with the Payroll monitoring group and generate a report containing information about those computers.
- Resources that appear in other reporting groups.
- A subset of resources within other reporting groups.

**Note:** Unlike monitoring groups, a resource can appear in one or more reporting groups.

You can create the following types of reporting group:

- Data Manager: **Computer, Filesystem**. Within a computer reporting group, you can include computer monitoring groups, specific computers, and other reporting groups. Within a file system reporting group, you can include file system monitoring groups, specific file systems, and other reporting groups.
- Data Manager for Databases: **Computer, Databases-Tablespaces**. Within a computer reporting group, you can include computer monitoring groups, specific computers, and other computer reporting groups. Within a databases-tablespaces reporting group, you can include databases-tablespace monitoring groups, specific databases or tablespaces, and other databases-tablespaces reporting groups.
- Disk Manager: **Storage Subsystems**. Within a storage subsystem reporting group, you can include storage subsystem monitoring groups, specific storage subsystems, and other storage subsystem reporting groups.

When you generate a report, storage statistics about a reporting group will appear in the corresponding "By Group" clause for that report. The following lists the reports that reflect reporting group information.

#### **Computer reporting groups (Data Manager and Data Manager for Databases )**

The following reports display computer reporting group information:

- Data Manager > Reporting > Capacity reports, Data Manager for Databases > Reporting > Capacity reports
- Data Manager > Reporting > Usage reports, Data Manager for Databases > Reporting > Usage reports
- Data Manager > Reporting > Usage Violation reports, Data Manager for Databases > Reporting > Usage Violation reports
- Data Manager > Reporting > Backup reports
- IBM Tivoli Storage Productivity Center > Rollup Reports > Capacity
- IBM Tivoli Storage Productivity Center > My Reports > Batch Reports > Create Batch Report > Report page, select one of the following:
  - Capacity > capacity\_report\_type > By Computer Group
  - Usage > usage\_reporting\_type > By Computer Group
  - Usage Violations > usage\_violation\_reporting\_type > By Computer Group
  - Backup > backup\_reporting\_type > By Computer Group
  - Rollup Reports > Capacity > capacity\_report\_type > By Computer Group
  - Rollup Reports > Database Capacity > rdbms\_type > By Computer Group

#### **Filesystem reporting groups (Data Manager)**

The following reports display filesystem reporting group information:

- Data Manager > Reporting > Capacity reports
- Data Manager > Reporting > Usage reports
- Data Manager > Reporting > Usage Violation > Constraint Violations reports
- Data Manager > Reporting > Backup reports
- IBM Tivoli Storage Productivity Center > My Reports > Batch Reports > Create Batch Report > Report page, select one of the following:
  - Capacity > capacity\_report\_type > By Filesystem Group
  - Usage > usage\_reporting\_type > By Filesystem Group
  - Usage Violations > Constraint Violations > By Filesystem Group
  - Backup > backup\_reporting\_type > By Filesystem Group

#### **Databases-Tablespaces reporting groups (Data Manager for Databases)**

The following reports display databases reporting group and tablespaces reporting group information:

- Data Manager for Databases > Reporting > Capacity reports
- Data Manager for Databases > Reporting > Usage reports
- Data Manager for Databases > Reporting > Usage Violations > Database-Tablespace Quotas reports

#### **Storage subsystem reporting groups**

The following reports display storage subsystem reporting group information:



- Data Manager > Reporting > Storage Subsystems > Storage Subsystem Views > By Storage System Group
- Data Manager > Reporting > Capacity > Disk Capacity > By Storage Subsystem Group
- IBM Tivoli Storage Productivity Center > My Reports > Batch Reports > Create Batch Report > Report page, select one of the following:
  - Storage Subsystems > Storage Subsystem Views > By Storage Subsystem Group
  - Capacity > Disk Capacity > By Storage Subsystem Group

## Creating a reporting group

To create a reporting group:

1. Expand **Reporting > Groups** under:
  - **Data Manager** to create computer and filesystem reporting groups.
  - **Data Manager for Databases** to create computer and databases-tablespaces reporting groups
  - **Disk Manager** to create storage subsystem reporting groups
2. Right click the type of group you want to create. A pop-up menu appears.
3. Select **Create <group type>** from the pop-up menu. The content pane displays a window for defining the group.
4. Define information about the group. See the IBM Tivoli Storage Productivity Center online Help for detailed information about the windows for creating a group.
5. Click **File > Save** to save the group definition. The group is saved to the database repository. You will now be able to generate and view reports based on these groups.

**Note:** You can define filters to select specific computers, filesystems, storage subsystems, tablespaces, and databases for a reporting group. Tivoli Storage Productivity Center uses these filters to dynamically include resources to the group as they are discovered within your environment. When defining dynamic filters for reporting groups, keep in mind that the filters you enter might or might not be case sensitive depending on the database under which the Tivoli Storage Productivity Center database repository is stored. For example:

- The dynamic filters you enter for reporting groups are case sensitive when the database repository is stored on Oracle, Sybase, and UDB/DB2.
- The dynamic filters you enter for reporting groups are not case sensitive when the database repository is stored on Microsoft SQL Server.

Keep in mind these case-sensitive issues when defining filters for reporting groups to prevent inaccurate reporting caused by unintended resources being included in reports.

## Editing a reporting group

To edit an existing reporting group:

1. Expand **Reporting > Groups** under:
  - **Data Manager** to edit computer and filesystem reporting groups.
  - **Data Manager for Databases** to edit computer and databases-tablespaces reporting groups
  - **Disk Manager** to edit storage subsystem reporting groups

2. Expand the type of the reporting group you want to edit. For example, expand **Disk Manager > Reporting > Groups > Storage Subsystem** to edit a storage subsystem reporting group.
3. Click the reporting group you want to edit. The content pane displays a window for editing the group.
4. Edit information about the group. See the IBM Tivoli Storage Productivity Center online Help for detailed information about the windows for creating a group.
5. Click **File > Save** to save the updated group definition.

**Note:** You can define filters to select specific computers, filesystems, storage subsystems, tablespaces, and databases for a reporting group. Tivoli Storage Productivity Center uses these filters to dynamically include resources to the group as they are discovered within your environment. When defining dynamic filters for reporting groups, keep in mind that the filters you enter might or might not be case sensitive depending on the database under which the Tivoli Storage Productivity Center database repository is stored. For example:

- The dynamic filters you enter for reporting groups are case sensitive when the database repository is stored on Oracle, Sybase, and UDB/DB2.
- The dynamic filters you enter for reporting groups are not case sensitive when the database repository is stored on Microsoft SQL Server.

Keep in mind these case-sensitive issues when defining filters for reporting groups to prevent inaccurate reporting caused by unintended resources being included in reports.

## Data Manager reports

Use Data Manager reports to view both overview and detailed information about your storage resources, including computers, disks, filesystems, storage subsystems, clusters, and users. These reports are constructed from the statistical information gathered by data collection jobs and accumulated in the database repository.

**Note:** You need to run appropriate data collection jobs on the storage resources before you generate reports. If you do not run the data collection jobs before generating reports, you might see inconsistent or inaccurate results.

### Data Manager - Report categories

Use the report categories within Data Manager to collect information about storage assets, availability, backup requirements, capacity, usage, usage violations, and monitored computer storage space.

The following table describes the categories of reports available in Data Manager.

| Report Category | Use these reports to do the following                                                                                                                                                                                                                                                                                | What job collects information for these reports?                                                    |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------|
| Asset reports   | View information about the agents, computers, disk controllers, hard disks, clusters, fabrics, storage subsystems, volumes, tape libraries, file systems, and general hardware inventory of storage assets within your environment.<br><br>See "Asset reports" on page 807 for more information about Asset reports. | Probes, discovery Data for System-wide Users and Monitored Directory reports are gathered by scans. |

| Report Category                | Use these reports to do the following                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | What job collects information for these reports?        |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------|
| Availability reports           | <p>View information about the availability of your storage, both from a network point of view and from a computer up-time perspective. This allows you to see what percentage of the time the storage in your enterprise (or on particular server or group of servers), is off-network due to network problems or perhaps is unavailable due to system downtime.</p> <p>See “Availability Reports” on page 831 for more information about Availability reports.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <p>Ping data: pings</p> <p>Computer uptime — probes</p> |
| TPC-wide Storage Space reports | <p>View storage capacity at a file system, computer, or entire network level. For a specific computer, group of computers, or all the computers in your enterprise, you can view the following:</p> <ul style="list-style-type: none"> <li>• How much storage you have</li> <li>• How much storage is being used</li> <li>• What percentage of storage is free for expansion</li> </ul> <p>See “TPC-wide Storage Space Reports” on page 832 for more information about Capacity reports.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Probes                                                  |
| Usage reports                  | <p>View information about the usage and growth of the consumption of your storage. Use these reports to do the following:</p> <ul style="list-style-type: none"> <li>• View which servers and file systems are experiencing the heaviest (or lightest) load of data and storage access.</li> <li>• Identify wasted space by pinpointing files that are no longer needed or have not been accessed for the longest time.</li> <li>• Associate users with the amount of storage they are consuming on specific file systems, groups of file systems, specific computers, groups of computers, and throughout the entire network.</li> <li>• Select the files against which to run an IBM Tivoli Storage Manager archive/backup job. Tivoli Storage Manager protects your organization's data from hardware failures and other errors by storing backup and archive copies of data on offline storage.</li> </ul> <p>See “Usage Reports” on page 892 for more information about Usage reports.</p> | Scans                                                   |

|                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                                                                                                  |
|---------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| <p><b>Report Category</b></p> <p><b>Usage Violation reports</b></p> | <p><b>Use these reports to do the following</b></p> <p>View information about the storage usage violations committed by users and user groups within your environment. Use these reports to do the following:</p> <ul style="list-style-type: none"> <li>• Create and enforce corporate storage usage policies. You can report on violations of quotas and constraints.</li> <li>• Select files in Constraint Violation reports against which to run an Tivoli Storage Manager archive or backup job.</li> <li>• Add files from Constraint Violation reports to an existing Tivoli Storage Manager archive or backup job.</li> </ul> <p>See “Usage Violation reports” on page 899 for more information about Usage Violation reports.</p> | <p><b>What job collects information for these reports?</b></p> <p>Scans, Quotas, Constraints</p> |
| <p><b>Backup reports</b></p>                                        | <p>Investigate the sizes of data that would be required to backup to achieve full backups, determine the amount of space required to perform an incremental backup of any or all of your data at any given point in time, and select files for archive/backup jobs.</p> <p>See “Backup reports” on page 901 for more information about Usage Violation reports.</p>                                                                                                                                                                                                                                                                                                                                                                       | <p>Scans</p>                                                                                     |
| <p><b>Monitored Computer Storage Space reports</b></p>              | <p>Use these host-based reports to view the data that is collected by probes for computers, Network Attached Storage, and hypervisors. These reports contain detailed information about the consumption of file system space and enable you to relate that space back to external or internal storage without counting shared space multiple times.</p> <p>See “Monitored Computer Storage Space reports” on page 903 for more information about Monitored Computer Storage Space reports.</p>                                                                                                                                                                                                                                            | <p>Probes</p>                                                                                    |

## Using By clauses

Use the By clauses associated with Data Manager reports as a method for viewing the data in those reports from different perspectives.

As you drill down through Data Manager reports, there is a set of options that repeats for each report type. Use these *By Clauses* to generate variations of the reports with different perspectives and groupings of data.

The available By Clauses are described in the following table:

| <b>By Clause</b>          | <b>Description</b>                                                                   |
|---------------------------|--------------------------------------------------------------------------------------|
| <b>By Directory</b>       | Use this clause to select the directories that you want to include in a report.      |
| <b>By Directory Group</b> | Use this clause to select the directory groups that you want to include in a report. |
| <b>By Disk</b>            | Use this clause to select the disks that you want to include in a report.            |

| <b>By Clause</b>            | <b>Description</b>                                                                                                                                                                                  |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>By Filesystem</b>        | Use this clause to select the file systems that you want to include in a report.                                                                                                                    |
| <b>By Filesystem Group</b>  | Use this clause to select the file system groups that you want to include in a report.                                                                                                              |
| <b>By Cluster</b>           | Use this clause to select the clusters or cluster resource groups that you want to include in a report. Cluster resource groups in a PowerHA or MSCS-clustered environment are listed as computers. |
| <b>By Computer</b>          | Use this clause to select the computers that you want to include in a report.                                                                                                                       |
| <b>By Computer Group</b>    | Use this clause to select the computer groups that you want to include in a report.                                                                                                                 |
| <b>By Storage Subsystem</b> | Use this clause to select the storage subsystems that you want to include in a report.                                                                                                              |
| <b>By Domain</b>            | Use this clause to select the Window Domains that you want to include in a report.                                                                                                                  |
| <b>Network Wide</b>         | The reports and charts that you see are on all data, network wide. Use this view to generate a global view of the storage and storage usage for your whole enterprise.                              |

## Asset reports

Use asset reports to view itemized statistics for your storage resources. Tivoli Storage Productivity Center provides a hierarchical view of your storage resources so that you can drill down to view your resources in greater detail.

The following is a description of each of the report types that are under **Data Manager > Reporting > Asset**.

| <b>Asset Reporting Type</b> | <b>Description</b>                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>By Cluster</b>           | View your storage assets based on PowerHA or MSCS cluster membership. You can drill down into a cluster and view its nodes and cluster resource groups. For each cluster resource group, you can drill down to show its controllers, disks, file systems, logical volumes, and exports or shares. For a cluster node, local resources can be viewed. |
| <b>By Computer</b>          | Drill down through your storage assets in a hierarchical manner. You can drill down into a computer and show its controllers, disks, file systems, logical volumes, and directories. You can then drill down further to view more detailed information such as the disks that are under a specific controller or file systems that are under a disk. |
| <b>By Hypervisor</b>        | Drill down through your hypervisors in a hierarchical manner. You can use this report to drill down and view details about the virtual machines, controllers, disks and file systems, or logical volumes.                                                                                                                                            |
| <b>By Storwize V7000U</b>   | Navigate through your Storwize V7000 Unified clusters and directories in a hierarchical manner. You can use this report to drill down and view details about the nodes, pools, Network Shared Disks, file systems, and exports that are associated with the clusters.                                                                                |

| Asset Reporting Type | Description                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| By OS Type           | View your storage assets based upon the operating systems that are running on those storage servers. This report type gives you an easy way to see what your OS dependencies are throughout your organization.<br><br>By drilling down on any or each of these different supported Operating System types, you can see how many of each server type you have and you can drill down to investigate the underlying storage. |
| By Storage Subsystem | View storage capacity at a computer, file system, storage system, volume, and disk level. These reports also show the relationships among the components of a storage system.                                                                                                                                                                                                                                              |
| System Wide          | View your storage assets as a whole, across the entire system. Use these reports to view lists of all of your disks, computers, file systems, Storage Resource agents and Data agents installed in your environment.                                                                                                                                                                                                       |

### By Cluster:

Use the By Cluster reports to drill down into a cluster to view its nodes and cluster resource groups.

- For a node, you can view its local resources, including controllers, disks, file systems, logical volumes, exports, shares, and directories.
  - For a cluster resource group, you can view controllers, disks, file systems, logical volumes, exports, and shares.
  - For a node, the number of processors that are reported are the logical processors. The amount of RAM that is reported is the available RAM.
1. Expand **Asset** → **By Cluster** node to display the nodes for each cluster that has been discovered and monitored by Data Manager.
  2. Click a node to view cluster information, or further expand the cluster node to drill down into greater detail.

### Cluster Information

Click the name of a cluster that appears under **Asset** → **By Cluster** to view information about that cluster such as:

- Name
- Domain
- IP address
- Discovered Time

### Controller Information

As you drill down into Asset Reporting By Cluster for a cluster node or cluster resource group, you can expand the assets under a controller, or you can select a controller to view its detailed information.

The name of the controller's computer is shown, as well as the driver name, instance number, bus number, address, and type of driver (for example, SCSI, IDE, etc).

**Note:** Host Bus Adapter (HBA) controllers must be connected to the SAN and see volumes in order to be displayed under the Controllers node.

### Disk Information

As you drill down into Asset Reporting By Cluster for a cluster node or cluster resource group, you can expand the assets under a disk, or you can select a disk to view its detailed information.

The disk detail information is shown in a tab dialog that is made of up three pages: **General**, **Paths**, and **Latest Probe**.

#### **General page**

Use the General page to view the computer name, the path name, SCSI Target ID, logical unit number, and the number of access paths. The disk information shown includes the hardware manufacturer of the disk, the model number of the disk, the firmware revision, the disk serial number, and the manufacture date along with the time when Data Manager discovered the disk. In the bottom area of the page is a breakdown of how the disk is allocated, including the disk allocation point and the length of the allocation.

#### **Latest Probe page**

Use the Latest Probe page to view the information that was obtained by Data Manager about this disk during the most recent probe of that disk. This information on this page includes:

- Sector size and sector information
- the number of heads
- the number of cylinders
- the logical block size
- the capacity of the disk
- the RPM or revolutions per minute that the disk spins at
- the elapsed time the disk has been powered on
- spare blocks
- the current write protection setting
- the cache enabled setting
- whether failure is predicted for this disk
- the time and date of the last Probe
- the unallocated space
- the number of primary defects
- the number of grown defects
- the NonMedium errors
- the recorded operating temperature

Some of these values are manufacturer specific, so if they are unavailable, then you will see **N/A** for a value. Additional information about the breakdown of recovered errors, unrecovered errors, and the number of bytes broken down by Write, Read, and Verify is shown in this report. Also included are three SCSI disk bit settings that you can turn on/off and the change will actually be made to the disk settings the next time a probe is run on that disk. These include on/off settings for:

- Automatic read reallocation
- Automatic write reallocation
- Read continuous (no error correction)

#### **Probe History page**

Use the Probe History page to view the history of probes that have been run on the disk you are reporting on. Each row on this report represents one time that a probe has been run against the disk. Each row shows what time the probe ran, the type of record (which could be daily, weekly or monthly), bytes written, bytes read, number of grown defects.

## Filesystem and Logical Volume Information

The Asset Reporting By Cluster Filesystem or Logical Volume information report contains two pages: **Filesystem** and **Logical Volume**.

### Filesystem page

Use this page to view a pie chart that depicts used space and free space on the file system and includes the following information:

- Computer name
- File system type
- Use count
- Mount point
- Physical size of the file system
- Capacity of the file system
- The time and date of the last probe and scan
- When the file system was discovered on the system by Data Manager
- If the file system has been removed then a removed date/time will be shown

### Logical Volume page

Use this page to view:

- The path name of the logical volume
- The capacity of the logical volume
- The use count
- The type of logical volume
- Whether it is a swap volume and mirrored
- A tree displaying the mapping of the logical volume on the disk area(s) that make up the logical volume

### Exports or Shares Information

The Asset Reporting By Computer Exports or Shares information report includes the following information:

- the name of the computer the share is on
- the export name that it will be known by
- the path on the computer the share resides upon
- the protocol being used and the date and time that Data Manager discovered the export or share

### Monitored Directory Information

To view a report on a directory that you are monitoring, expand the **Monitored Directories** node on the navigation tree and select the directory on which you want to report. This information is available for clustered nodes, but not cluster resource groups.

The Asset Reporting By Cluster Monitored Directory report includes the following information:

- computer name the directory resides on
- name of the directory
- owner
- OS Group
- total size of the directory
- number of files in the directory
- count of subdirectories in the directory
- average file size
- access time



- modification time and creation time of the directory
- last scan time and the time that Data Manager discovered the directory
- file attributes that are set
- list of the Directory Groups that the directory is a part of

### **By Computer:**

Use the By Computer reports to drill down into a computer and view detailed information about its controllers, disks, file systems, logical volumes, and directories.

### **Computer Information**

As you drill down into Asset Reporting By Computer, you can expand the assets under a computer, or you can select a computer to view the details for that computer. The detailed report on the right shows the following information:

- Computer
- Host ID: the identifier of the host computer in hexadecimal format. Currently, IBM Tivoli Storage Productivity Center collects Host ID information for Solaris computers only. All other platforms display GUID information in the Host ID column.
- Group and Domain
- Network Address, IP address
- Time zone
- Hardware manufacturer and model numbers
- Processor type
- Count and speed
- Amount of RAM
- Operating System (OS) information such as:
  - Type of OS and version
  - The storage attributes of the server
  - The computer disk capacity, unallocated space, and free space

### **Controller Information**

As you drill down into Asset Reporting By Computer, you can expand the assets under a controller, or you can select a controller to view the details for that controller. The name of the computer controller is shown, as well as the driver name, instance number, bus number, address, and type of driver (for example, SCSI, IDE, and so on).

**Note:** Host Bus Adapter (HBA) controllers must be connected to the SAN and see volumes in order to be displayed under the Controllers subnode.

### **Disk Information**

As you drill down into Asset Reporting By Computer, you can expand the assets under a disk, or you can select a disk to view the details for that disk. The disk detail information is shown in a tab dialog that is made of up four tabs: **General**, **Paths**, **Latest Probe**, and **Probe History**.

#### **General page**

The General page includes the computer name, the path name, SCSI Target ID, logical unit number, the number of access paths.

The disk information shown includes:

- The hardware manufacturer of the disk
- The model number of the disk

- The disk firmware revision
- The disk serial number
- The disk manufacture date
- Indicators that specify whether the disk is encrypted or encryptable
- The time when Data Manager discovered the disk

In the bottom area of the tab is a breakdown of how the disk is allocated, including the disk allocation point and the length of the allocation.

**Paths** The **Paths** page shows the following information:

- Host
- OS type
- Path
- Controller
- Instance
- Bus number
- SCSI target id
- Logical unit number

#### **Latest Probe page**

The **Latest Probe** page shows the information that was obtained by Data Manager about this disk during the most recent probe of that disk. This information on this tab includes:

- Sector size and sector information
- The number of heads
- The number of cylinders
- The logical block size
- The capacity of the disk
- The RPM or revolutions per minute that the disk spins at
- The elapsed time the disk has been powered on
- Spare blocks
- The current write protection setting
- The cache enabled setting
- Whether failure is predicted for this disk
- The time and date of the last Probe
- The unallocated space
- The number of primary defects
- The number of grown defects
- The NonMedium errors
- The recorded operating temperature

Some of these values are manufacturer-specific. If the values are unavailable, **N/A** is displayed.

Additional information about the breakdown of recovered errors, unrecovered errors, and the number of bytes broken down by Write, Read, and Verify is shown in this report.

Also included are three SCSI disk bit settings that you can turn on/off and the change will actually be made to the disk settings the next time a probe is run on that disk. These include on/off settings for:

- Automatic read reallocation
- Automatic write reallocation
- Read continuous (no error correction)

### **Probe History page**

The **Probe History** page shows the history of probes that have been run on the disk you are reporting on. Each row on this report represents one time that a probe has been run against the disk. Each row shows what time the probe ran, the type of record (which could be daily, weekly, or monthly), bytes written, bytes read, number of grown defects.

### **Volume Group Information**

Data Manager provides support for the following volume managers:

- Veritas Volume Manager (Solaris)
- Logical Volume Manager (AIX)

These volume managers allow you to group logical volumes and disks. As you drill down into Asset Reporting By Computer for an AIX or Solaris computer, you can expand the assets under a volume group to see the group disks and file systems or logical volumes.

You can also select a volume group to view detailed information about it. The group name is shown, as well as the type (LVM or Veritas), capacity, free space, number of volumes, number of disks, discovered time, and probe time.

### **File System / Logical Volume Information**

The Asset Reporting By Computer File System / Logical Volume information report contains two pages: **File System** and **Logical Volume**.

#### **File System page**

This page shows a pie chart depicting used space and free space on the file system and includes the following information:

- Computer name
- File system type
- Use count
- Mount point
- Physical size of the file system
- Capacity of the file system
- The time and date of the last probe and scan
- When the file system was discovered on the system by Data Manager
- If the file system has been removed, then a removed date/time is shown

Also shown for UNIX-based systems are:

- The maximum file count
- Used inodes
- Free inodes. Inodes are areas reserved for index entries within the UNIX file systems. For Windows computers, these values are zero.
- The amount of used space and free space is shown

#### **Logical Volume page**

The information on this page includes:

- The path name of the logical volume
- The capacity of the logical volume
- The use count
- The type of logical volume
- Whether it is a swap volume and mirrored

- A tree displaying the mapping of the logical volume on the disk areas that make up the logical volume

### **Exports/Shares Information**

The Asset Reporting By Computer Exports/Shares information report includes the following information:

- The name of the computer the share is on
- The export name that it will be known by
- The path on the computer the share resides upon
- The protocol being used and the date/time that Data Manager discovered the export/share

### **Monitored Directory Information**

To view a report on a directory that you are monitoring, expand the **Monitored Directories** node on the navigation tree and select the directory on which you want to report.

**Note:** This information is not shown for cluster resource groups. The Asset Reporting By Computer Monitored Directory report includes the following information:

- Computer name the directory resides on
- Name of the directory
- Owner
- OS Group
- Total size of the directory
- Number of files in the directory
- Count of subdirectories in the directory
- Average file size
- Access time
- Modification time and creation time of the directory
- Last scan time and the time that Data Manager discovered the directory
- File attributes that are set
- List of the Directory Groups that the directory is a part of

The Asset Reporting By Computer Monitored Directory report includes the following information:

- Computer name the directory resides on
- Name of the directory
- Owner
- OS Group
- Totalsize of the directory
- Number of files in the directory
- Count of subdirectories in the directory
- Average file size
- Access time
- Modification time and creation time of the directory
- Last scan time and the time that Data Manager discovered the directory
- File attributes that are set
- List of the Directory Groups that the directory is a part of

### **By OS Type:**

Use the By OS Type Asset report to view your storage assets based upon the operating systems under which they are running. This provides you with an easy way to see what your OS dependencies are throughout your organization. You can drill down into any of the storage resources that appear under an operating system node to access all of the same sub-reports that as described for other asset reports.

## By Storage Subsystem:

Browse your storage subsystem assets in a hierarchical manner. Data Manager provides reports for individual storage subsystems and any disk array subsystems whose SMI-S Providers are CTP certified by SNIA for SMI-S 1.0.2 or later.

For disk array systems, you can view information about:

- Disk groups
- Array sites
- Ranks
- Storage pools
- Disks
- Volumes

For virtualization systems, you can view information about:

- SAN Volume Controller clusters
- Storwize V7000
- Storwize V7000 Unified
- Managed disk groups
- Managed disks
- Virtual disks
- Virtual disk copies

Keep in mind the following considerations:

- A volume name in a report that is marked with an asterisk (\*) indicates that the volume is a space-efficient volume. The \* is not part of the actual volume name.
- The total capacity value for storage pools in TotalStorage Enterprise Storage Server subsystems includes only the formatted (usable) capacity for storage pools.
- For the following storage systems, you can drill down from storage subsystems and show managed disks, managed disk groups, and virtual disks:
  - SAN Volume Controller
  - Storwize V7000
  - Storwize V7000 Unified

For Storwize V7000 and Storwize V7000 Unified, you can also drill down from the internal disks node to view the internal storage on the system.

- Information that is displayed for managed disks and internal disks is the same as for storage subsystem disks. See “By Storage Subsystem” for a description of the available views.
- Information that is displayed for virtual disks is the same as for storage subsystem volumes. See “By Storage Subsystem” for a description of the available views.
- For IBM FAStT, HP StorageWorks, and the XIV system, you can drill down into storage subsystems and show storage pools, disks, and volumes.

**Note:** The total capacity value for storage pools in these storage subsystems includes raw (unconfigured) capacity—it does not include formatted/usable capacity.

- For Hitachi Data Systems TagmaStore subsystem, the relationship between a disk or volume is not displayed if the volume is using imported extents (virtual disks).

**Note:** The total capacity value for storage pools in HDS subsystems includes the formatted/usable capacity for storage pools only.

- Some of the columns in the totals row for these reports can count storage multiple times if the report includes the following:
  - A storage virtualizer like a SAN Volume Controller
  - The disk arrays that host the backend storage to a storage virtualizer like a SAN Volume Controller

In this situation, disks and logical unit numbers (LUNs) appear to belong to both a storage virtualizer like SAN Volume Controller and the disk arrays, and are counted multiple times in some of the columns in the totals row. To avoid this situation, filter the report to display only the storage virtualizer or the disk arrays. If this situation is unavoidable, then some of the columns in the totals row can be adjusted as follows:

1. Use the **Physical Disk Space** column to determine the actual physical disk space on all arrays. This disk space is not counted multiple times.
2. Sum the **Disk Space** column for the disk arrays and subtract this result from the totals row to determine the total backend storage for a virtualizer.
3. To prevent counting the backend volumes multiple times, subtract the **Backend Volume Space** column from the **Volume Space** column in the totals row to determine the amount of volume space.

See “Report considerations and limitations” on page 781 for additional information about reports.

### Disk Groups

When you expand the **Disk Group** node, the navigation tree shows a list of disk groups for the storage subsystem. You can expand the name of a disk group to view the disks and volume spaces within that disk group. Drill down through both Volume Spaces and Disks subnodes to view:

- the disks and volumes associated with a volume space
- the volumes with segments on a disk
- the disks containing segments from a volume
- detailed information about a disk

### Array Sites

(IBM Tivoli Storage DS8000 only) When you expand the **Array Site** node, the navigation tree shows a list of array sites for the storage subsystem. Click **Array Site** to view a description of the information displayed about array sites. You can expand the name of an array site to view the ranks, storage pools, and disks associated with that array site.

### Storage Pools

When you expand the **Storage Pool** node, the navigation tree shows a list of storage pools for the storage subsystem. You can expand the name of a storage pool to view the disks and volumes within that storage pool, including:

- the volumes with segments on a disk
- the disks containing segments from a volume
- detailed information about a disk

**Disks** When you expand the **Disks** node, the navigation tree shows a list of disks

for the storage subsystem. You can expand the name of disk to view the volumes associated with that disk. Click the name of a volume associated with a disk to view detailed information about that volume.

### **Volume Spaces**

When you expand the **Volume Spaces** node, the navigation tree shows a list of volume spaces for the storage subsystem. You can expand the name of a volume space to view the disks and volumes within that volume space. Drill down through both the Disks and volumes subnodes to view:

- the volumes with segments on a disk
- the disks containing segments from a volume
- detailed information about a disk
- detailed information about a volume

### **Volumes**

When you expand the volumes node, the navigation tree shows a list of volumes for the storage subsystem. Click the name of a volume to view detailed information about that volume.

#### **Note:**

- Volume Reporting is not available for IBM FAStT FlashCopy logical drive and repository logical drive.
- After creating a volume on a storage subsystem, you must run a probe on that storage subsystem to collect accurate information about disk allocation.

### **Managed Disk Groups**

When you expand the Managed Disk Group node, the navigation tree shows a list of disk groups that are managed by the following storage systems:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

You can expand the name of a disk group to view the physical disks and virtual disks that are associated with it.

### **Managed Disks**

When you expand the Managed Disk node, the navigation tree shows a list of disks that are managed by the following storage systems:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

You can expand the name of a disk to view the virtual disks and any available internal disks that are associated with it.

### **Virtual Disks and Virtual Disk Copies**

There are three nodes in the navigation tree that are related to virtual disks (VDisks). These nodes might be included under the nodes for the SAN Volume Controllers in your environment:

- **Virtual Disks** (at the same level as **Managed Disk** and **Managed Disk Groups**): Expand this node to view a list of the logical disks that a cluster provides to the SAN. You can further expand each virtual disk (VDisk) to view its mirror copies (primary copies and additional copies).

- **Virtual Disks** (under **Managed Disk** and **Managed Disk Groups**): Expand this node to view the VDisks for which the primary copies are on the selected Managed Disk.
- **Virtual Disk Copies** (under **Managed Disk** and **Managed Disk Groups**): Expand this node to view all the virtual disk copies that are on the selected Managed Disk. The primary copies are not listed under this node because they are displayed under the **Virtual Disks** node.

**Note:**

- The **Virtual Disk Copies** node is displayed in the navigation tree only if a monitored SAN Volume Controller supports virtual disk mirrors. If it supports virtual disk mirrors, all virtual disks will have a related primary virtual disk copy regardless of whether the virtual disk is mirrored. This is because all virtual disks on a SAN Volume Controller that supports virtual disk mirrors are primary copies that can be updated to have mirrored disks or virtual disk copies.
- The icons for all virtual disk copies are shown as folders in the navigation tree because they represent logical extensions of a virtual disk.
- If the virtual disk is a virtual disk copy, the primary virtual disk copy is noted by -0 in the name of the logical volume, and all secondary copies are noted by a hyphen and a value larger than zero. If the virtual disk is not a virtual disk copy, then only one logical volume is created and is noted by -0. The logical disk representation of a virtual disk should always exist, but an probe is started for SAN Volume Controller 4.3.1 or later, the logical disk representation of the virtual disk expands to one or more logical volumes where virtual disk copies are represented by multiple logical volumes, and non-virtual disk copies are represented by a single logical volume.

**Internal Disks**

Internal storage information about the Storwize V7000 or Storwize V7000 Unified systems in your environment is displayed under the **Internal Disks** nodes in the navigation tree:

**Internal Disks (at the same level as Virtual Disks, Managed Disk, and Managed Disk Groups)**

Expand this node to view a list of the internal disks that a Storwize V7000 or Storwize V7000 Unified system provides to the SAN.

**Internal Disks (under Managed Disk Groups > Managed Disk)**

The internal disks for this node are only displayed if there are internal disks that are available on the managed disk.

**Internal Disks (under Managed Disk)**

The internal disks for this node are only displayed if there are internal disks available.

**System-wide:**

Use Asset > System-wide reports to perform the following tasks:

- View your storage assets as a whole across the entire system. You can view lists of all of your disks, computers, storage subsystems, disk/volume groups, disks, file systems, logical volumes, storage subsystem volumes, exports or shares, monitored directories, users, and OS users. Additionally, you can view aggregated reports based upon certain filters, such as grown defects per disk.



- View detailed information about Storage Resource agents and Data agents. Use these reports to monitor your agents and track their status.
- The **Selection** button is disabled for system-wide reports because all the monitored storage entities are automatically displayed for each report type. For example, all the computers monitored by IBM Tivoli Storage Productivity Center are displayed in the **System-wide > Computers** report.

**Note:** A volume name in a report that is marked with an asterisk (\*) indicates that the volume is a space-efficient volume. The \* is not part of the actual volume name. See “Report considerations and limitations” on page 781 for additional information about reports.

System-wide reports include options to drill down into storage resources for more detail, view charts of current statistics, and generate historical charts to show storage trends. The following icons are available on some system-wide reports:

**Chart icon**  Chart icon

Click this button to generate charts for the selected rows in the table. You can generate a chart based on the statistics currently displayed in the report, or you can generate a history chart that provides trending data for the selected storage subsystems. See Viewing historical trends for more information about history charts.

**Magnifying glass icon**  Magnifying glass icon

Click this button to view detailed information about a specific file system or logical volume in a row.

**Agents reports**

Use these reports to track and monitor the Storage Resource agents and Data agents that are deployed in your environment. With this report, you can export a comma-separated file containing the name and status of an agent. Specifically, use this report to perform the following tasks:

- View information about the agents deployed in your environment
- Perform actions against one or more agents shown in the report

Each agent that is shown in the report has a magnifying glass icon next to it. Click this icon to display detailed information about an agent (which is identical to the information on the **Administrative Services > Data Sources > Data/Storage Resource agents** node in the navigation tree).

You can also select one or more rows on the report and right-click to display a menu. Use the options in this menu to perform an action against multiple agents at the same time. The menu contains all the other operations available from the **Administrative Services > Data Sources > Data/Storage Resource agents** node in the navigation tree.

**Computers reports**

Use these reports to view all the computers in the system, sorted in order of the time they were last booted, probed, or discovered. This information is valuable in environments where computers become more unreliable after they have been up and running for long periods of time.

**Storage Subsystem reports**

Use these reports to view all your monitored storage subsystems and SAN Volume Controllers, sorted by their total capacity (By Capacity) or alias (By Storage Subsystem).

### **Unmanaged Virtual Machines reports**

Use these reports to view detailed information about the unmanaged virtual machines in your environment. Each row in the report gives you information for each unmanaged virtual machine.

### **Virtual Machines With No VMware reports**

Use these reports to view detailed information about the virtual machines that have Storage Resource agents or Data agents installed but are not discovered through the VMWare data sources. Each row in the report gives you information for each virtual machine.

### **Disk/Volume Groups reports**

Data Manager provides support for the following volume managers:

- Veritas Volume Manager (Solaris)
- Logical Volume Manager (AIX)

These volume managers enable you to create groups of logical volumes and disks. Use these reports to view all your disk and volume groups, sorted by their computer (By Computer), storage capacity (By Capacity) total freespace (By Freespace), time when they were last probed (By Probe Time), and time when they were discovered (By Discovered Time).

### **Disks reports**

Use these reports to view all the disks in the system, sorted in order of the amount of unallocated space they have on them (By Unallocated Space), how many recovered and unrecovered errors they have, time when they were probed (By Probe Time), and time when they were discovered (By Discovered Time).

Use the **Asset Reporting → System-wide → Disks by Grown Defects** report to view the following information:

- All the disks in the system sorted in order of the number of grown defects they have on them
- What disks in your system are experiencing problems and might be in danger of an imminent failure

The **Asset Reporting → System-Wide → Disks by Primary Defects** report returns no data if disk defects are not detected. Use this report to view the following information:

- All the disks in the system sorted in order of the number of primary defects they have on them
- What disks in your system are in good shape and which ones are more or less defective

### **File systems / Logical Volumes reports**

Use these reports to view detailed information about file systems and logical volumes on your monitored computers. A logical volume is a partition on a disk that does not have a file system (raw partition). Use these reports to view the following information:

- Available file system space on monitored computers.
- All the file systems in the system that is sorted in order of the amount of free space they have on them.
- Which file systems have less free space and might be in danger of running out of free space.
- Which file systems have too much free space and where space consolidation might be necessary.

- A list of virtual disk files which are not assigned to any of the virtual machines. This situation represents wasted storage capacity.
- The path from the file system on a virtual machine to the backend storage volumes on a hypervisor or mapped to the hypervisor.

**Note:** The values in the By Available File System Space report are based on data that is collected by probes. Probes collect space information as measured in blocks (1 block = 512 bytes). Tivoli Storage Productivity Center calculates file system space, including used space and free space, based on the number of blocks it retrieves from a probe. It multiplies the number of blocks by 512 and displays the resulting value in the report.

### **Volumes reports**

Use these reports to view the following information:

- Information about the storage subsystems that host volumes
- Storage capacity of the volumes hosted by storage subsystems
- Information about the storage subsystems that host volumes sorted by the date/time when they were probed

### **Exports or Shares report**

Use this report to view all the Exports or Shares in your environment, sorted in the order that they were discovered.

### **Monitored Directories reports**

Use these reports to view all the monitored directories in the system, sorted in order of the amount of storage space they are taking up (By Space Used), the time when they were last scanned (By Scan Time), the time when they were discovered (By Discovered Time), and when they were removed (Removed Directories).

### **Unmanaged Computers report**

Use this report to view all of the new computers that have been introduced into your system and that do not have Storage Resource agents or Data agents installed on them. It also lists other computers that you have not put under Data Manager control.

### **Users report**

Use this report to view all the users in your environment sorted in the order that they were discovered.

### **OS User Groups report**

Use this report to view all the OS User Groups in your environment sorted in the order when they were discovered.

### **By Storwize V7000U:**

Use By Storwize V7000U reports to drill down into the components of a Storwize V7000 Unified storage system. These components include the clusters, nodes, pools, Network Shared Disks, file systems, and exports that are associated with the storage system.

To view Storwize V7000 Unified asset reports, follow these steps:

1. Expand **Data Manager > Reporting > Asset > By Storwize V7000U**. The Storwize V7000 Unified storage systems that were discovered by Tivoli Storage Productivity Center are displayed.
2. Expand a storage system to drill down for greater detail. The following tabbed pages are displayed:

- **Storage Subsystem Details:** Use this page to view detailed information about the Storwize V7000 Unified storage system, including the status, vendor, type and model, and the available and used space. You can also use this page to modify some information about the Storwize V7000 Unified storage system.
- **File Module Details:** Use this page to view information about Storwize V7000 File Module.

You can drill down to view information about the components of the Storwize V7000 Unified storage system. When you click a component, report pages for that component are displayed as described in the following sections.

*Storwize V7000U - Cluster Information:*

This page shows general information about a selected cluster that is associated with a Storwize V7000 Unified storage system.

The following information is displayed on this page.

| Field Name | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cluster    | The name of the cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Cluster ID | The identifier of the cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Type       | <p>The cluster can be one or more of the following types:</p> <p><b>Interface</b><br/>An interface cluster contains nodes that connect a system to an IP network for file-serving capabilities by using service protocols.</p> <p><b>Management</b><br/>A management cluster contains nodes that are used to configure, administer, and monitor a system.</p> <p><b>Storage</b><br/>A storage cluster contains nodes that provide the storage and file system for the data in a system.</p> |

*Storwize V7000U - Node Information:*

This page shows general information about a selected node that is associated with a Storwize V7000 Unified storage system.

The following information is displayed on this page.

| Field Name  | Description                  |
|-------------|------------------------------|
| Node        | The name of the node.        |
| Description | The description of the node. |
| IP Address  | The IP address for the node. |

| Field Name | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Role       | <p>The role of the node. The node can serve one or more of the following roles:</p> <p><b>Interface</b><br/>An interface node connects a system to an Internet Protocol (IP) network for file-serving capabilities by using service protocols.</p> <p><b>Management</b><br/>A management node is used to configure, administer, and monitor a system.</p> <p><b>Storage</b><br/>A storage node provides the storage and file system for the data in a system.</p> |
| Status     | <p>The status of the node. The status can be one of the following values:</p> <p><b>Normal</b><br/>The node is operating normally.</p> <p><b>Critical</b> The node is either not operating or has problems.</p> <p><b>Warning</b><br/>At least one component of the node is not operating or has problems.</p> <p><b>Unreachable</b><br/>Tivoli Storage Productivity Center cannot reach the node.</p>                                                            |

*Storwize V7000U - Pool Information:*

This page shows general information about a selected pool that is associated with a Storwize V7000 Unified storage system.

The following information is displayed on this page.

| Field Name      | Description                                     |
|-----------------|-------------------------------------------------|
| Pool            | The name of the pool.                           |
| Capacity        | The capacity of the pool.                       |
| Available Space | The amount of unused space that is in the pool. |

*Storwize V7000U - NSD Information:*

This page shows general information about a selected Network Shared Disk (NSD) that is associated with a Storwize V7000 Unified storage system.

The following information is displayed on this page.

| Field Name | Description          |
|------------|----------------------|
| NSD        | The name of the NSD. |

| Field Name      | Description                                                                                                                                                                                                                                                                                                                                                          |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NSD ID          | The name of the cluster that the NSD is associated with.                                                                                                                                                                                                                                                                                                             |
| Failover Group  | The failure group that the NSD belongs to.                                                                                                                                                                                                                                                                                                                           |
| Type            | The NSD type. The possible values are: <ul style="list-style-type: none"> <li>• Metadata and Data</li> <li>• Metadata Only</li> <li>• Data Only</li> </ul>                                                                                                                                                                                                           |
| Pool            | The pool that the NSD belongs to.                                                                                                                                                                                                                                                                                                                                    |
| Status          | The status of the NSD. The possible values are: <p><b>Normal</b><br/>The NSD is operating normally.</p> <p><b>Critical</b> The NSD is either not operating or has problems.</p> <p><b>Warning</b><br/>At least one component of the NSD is not operating or has problems.</p> <p><b>Unreachable</b><br/>Tivoli Storage Productivity Center cannot reach the NSD.</p> |
| Disk Space      | The total disk space for the NSD.                                                                                                                                                                                                                                                                                                                                    |
| Available Space | The amount of unused disk space for the NSD.                                                                                                                                                                                                                                                                                                                         |

*Storwize V7000U - NSD Probe History:*

This page shows the probe history for a selected Network Shared Disk (NSD) that is associated with a Storwize V7000 Unified storage system.

The following information is displayed on this page.

| Field Name             | Description                                                                                                                                                   |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Time of Probe          | The time that the probe ran.                                                                                                                                  |
| Record Type            | The record type that is shown in the report. Valid values are: <ul style="list-style-type: none"> <li>• Daily</li> <li>• Weekly</li> <li>• Monthly</li> </ul> |
| Bytes Written          | The total number of bytes that were written since the NSD was formatted or manufactured.                                                                      |
| Bytes Read             | The total number of bytes that were read since the NSD was formatted or manufactured.                                                                         |
| Grown Defects          | The number of defects that were detected on the NSD after the NSD was used.                                                                                   |
| Recovered Write Errors | The number of write errors that were successful on a subsequent attempt.                                                                                      |

| Field Name                | Description                                                                                                                                      |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| Unrecovered Write Errors  | The number of write errors that were not successful on a subsequent attempt.                                                                     |
| Recovered Read Errors     | The number of read errors that were successful on a subsequent attempt.                                                                          |
| Unrecovered Read Errors   | The number of read errors that were not successful on a subsequent attempt.                                                                      |
| Recovered Verify Errors   | The number of verify errors that were successful on a subsequent attempt.                                                                        |
| Unrecovered Verify Errors | The number of verify errors that were not successful on a subsequent attempt.                                                                    |
| Non-Medium Errors         | The number of errors that are not related to the physical disk platter, such as mechanical problem in the head movement or communication errors. |
| Available Disk Space      | The amount of unused space on the NSD.                                                                                                           |
| Primary Defects           | The number of defects on the NSD when the NSD was new.                                                                                           |
| Temperature               | The temperature of the NSD in degrees Celsius.                                                                                                   |

*Storwize V7000U - Correlated Storage Subsystem Volume:*

This page shows information about the relationship between a selected Network Shared Disk (NSD) and the associated virtual disk (volume) that is on the Storwize V7000 Unified storage system.

The following information is displayed on this page.

| Field Name                     | Description                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Virtual Disk Name              | The name of the virtual disk.                                                                                                                                                                                                                                                                                                                                                         |
| Unavailable Virtual Disk Space | Amount of storage space for Redundant Array of Independent Disks (RAID) or mirror redundancy.                                                                                                                                                                                                                                                                                         |
| Type                           | Organization of the virtual disk. Values are: <ul style="list-style-type: none"> <li>• Sequential</li> <li>• Striped</li> </ul>                                                                                                                                                                                                                                                       |
| RAID Level                     | The RAID level of the virtual disk. For example, 0, 1, 5, and so on.                                                                                                                                                                                                                                                                                                                  |
| Is Encryptable                 | Indicates whether the virtual disk is encryptable. Encryptable means that the virtual disk uses full disk encryption, but does not belong to a defined System Storage DS8000 or System Storage DS5000 encryption group. Values are: <ul style="list-style-type: none"> <li>• True (the virtual disk is encryptable)</li> <li>• False (the virtual disk is not encryptable)</li> </ul> |

| Field Name       | Description                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Is Encrypted     | Indicates whether the virtual disk is encrypted. Encrypted means that the volume uses full disk encryption and belongs to a defined System Storage DS8000 or System Storage DS5000 encryption group. Values are: <ul style="list-style-type: none"> <li>• True (the virtual disk is encrypted)</li> <li>• False (the virtual is not encrypted)</li> </ul>                                                 |
| Easy Tier        | Indicates whether Easy Tier is on or off for the virtual disk.                                                                                                                                                                                                                                                                                                                                            |
| Easy Tier Status | Indicates the Easy Tier functions that are active for the virtual disk copy. Values are: <ul style="list-style-type: none"> <li>• active (can move extents of this virtual disk copy for performance [automatic data placement])</li> <li>• inactive (no Easy Tier function is active)</li> <li>• measured (statistics are being gathered for the virtual disk copy, but no extents are moved)</li> </ul> |
| Tiers            | The tier levels that are used by the virtual disk. Each tier is displayed as a row in a table.                                                                                                                                                                                                                                                                                                            |
| Flash Copy       | Indicates whether the virtual disk or copy is in a FlashCopy (FC) relationship and whether it is an FC source or FC target. A virtual disk or copy that is not in an FC relationship is displayed with a None value regardless of whether the virtual disk is on a system that supports FC.                                                                                                               |
| Attributes       | Indicates whether the disk is surfaced or not surfaced.                                                                                                                                                                                                                                                                                                                                                   |

*Storwize V7000U - File System Information:*

This page shows general information about a selected file system that is associated with a Storwize V7000 Unified storage system.

The following information is displayed on this page.

| Field Name        | Description                                                                                             |
|-------------------|---------------------------------------------------------------------------------------------------------|
| Computer          | The name of the computer that contains the file system.                                                 |
| File System Type  | The file type, which is General Parallel File System (GPFS).                                            |
| Mount Point       | The name or mount point (UNIX/Linux) of the file system (for example, c:\, d:\, /opt, or /export/home). |
| Physical Space    | The physical size of the file system.                                                                   |
| File System Space | The capacity of the file system.                                                                        |
| Probe Last Run    | The date and time that the last probe ran for the file system.                                          |



| Field Name         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Scan Last Run      | The time that the last scan ran for the file system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Discovered Time    | The date and time that the file system was discovered.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Removed Time       | The date and time that the file system was removed, if applicable.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Maximum File Count | <p>The total number of files within a file system. This value does not include files on file systems that were not scanned.</p> <p>On Windows systems, AIX systems, and systems that are using General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for this file. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.</p>                                                                                                                                                               |
| Used I-Nodes       | The number of used i-nodes in the file system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Free I-Nodes       | The number of free i-nodes within a file system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Used Space         | <p>Amount of used storage space within a file system.</p> <p>On some UNIX file systems, this value does not include storage space that is reserved by the operating system. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value:</p> <ul style="list-style-type: none"> <li>• File systems that become missing after a probe.</li> <li>• File systems that are mounted remotely and are not identified to a network-attached storage (NAS) server.</li> <li>• File systems that are not mounted.</li> </ul> <p>The amount of used and free space is also shown in a chart at the bottom of the page.</p> |
| Free Space         | The amount of unused storage space in the file system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

*Storwize V7000U - File Set Information:*

This page shows general information about a selected file set that is associated with a Storwize V7000 Unified storage system.

The following information is displayed on this page.

| Field Name | Description               |
|------------|---------------------------|
| File Set   | The name of the file set. |

| Field Name   | Description                                                                                                                                                                                                                                                                                   |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Path         | The path for the file set. The path is displayed only if Linked is displayed in the <b>Status</b> field.                                                                                                                                                                                      |
| File System  | The name of the file system that is associated with the file set.                                                                                                                                                                                                                             |
| Status       | The status of the file system. Valid values are:<br><br><b>Linked</b> The file set is linked to the file system that is displayed in the <b>File System</b> field.<br><br><b>Unlinked</b><br>The file set is not linked to the file system that is displayed in the <b>File System</b> field. |
| Used Space   | The amount of space that is used by the file set.                                                                                                                                                                                                                                             |
| Used I-Nodes | The number of i-nodes that are used by the file set.                                                                                                                                                                                                                                          |

*Storwize V7000U - Export Information:*

This page shows general information about a selected export that is associated with a Storwize V7000 Unified storage system.

The following information is displayed on this page.

| Field Name      | Description                                                                                                                                                                                                                                                                                                                                         |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Export or Share | The name of the computer that the export is on.                                                                                                                                                                                                                                                                                                     |
| Cluster         | The name of the cluster that the export is associated with.                                                                                                                                                                                                                                                                                         |
| Path            | The path on the computer where the export is located.                                                                                                                                                                                                                                                                                               |
| Discovered      | The date and time that the export was discovered.                                                                                                                                                                                                                                                                                                   |
| Active          | Indicates whether the export is active or inactive within the cluster. Valid values are: <ul style="list-style-type: none"> <li>• Yes: The export is active.</li> <li>• No: The export is inactive. If an export is inactive, it is included in the list of exports, but the data that is associated with the export cannot be accessed.</li> </ul> |

*Storwize V7000U - Configured Protocol Information:*

This page shows the protocols for a selected export that is associated with a Storwize V7000 Unified storage system.

Protocols are specified for an export when the export is created. For each protocol that is listed on this page, either Configured or Not Configured is shown to

designate whether the protocol as been specified for the export. In addition, the information in the following tables is provided for the Common Internet File System (CIFS) and Network File System (NFS) protocols.

### Protocol Type CIFS

| Field Name                       | Description                                                                                                                                                                                                                                                                               |
|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Read Only                        | Specifies whether the export is read only. If the export is read only, files cannot be modified or created on this export. Values are: <ul style="list-style-type: none"> <li>• Yes (files cannot be modified or created)</li> <li>• No (files can be modified or created)</li> </ul>     |
| Browsable by Client              | Specifies whether the export is hidden from the client. The export is not shown in the list of available exports. Values are: <ul style="list-style-type: none"> <li>• Yes (The export is not hidden from the client.)</li> <li>• No (The export is hidden from the client.)</li> </ul>   |
| Hide Objects with No Read Access | Specifies whether the export is excluded from lists of files and directories when the user does not have permission to those files and directories. Values are: <ul style="list-style-type: none"> <li>• Yes (The export is excluded.)</li> <li>• No (The export is included.)</li> </ul> |
| Comment                          | A user-defined comment for the CIFS export. If a comment is not available, this field is blank.                                                                                                                                                                                           |
| Client/Group/SID                 | The client, group, or security identifier (SID) that can access the CIFS protocol.                                                                                                                                                                                                        |
| Access                           | Specifies whether the client, group, or SID has access to the export. <ul style="list-style-type: none"> <li>• Allowed</li> <li>• Denied</li> </ul>                                                                                                                                       |
| Value                            | Specifies the level of access that is granted to the client, group, or SID. Values are: <ul style="list-style-type: none"> <li>• Full</li> <li>• Read</li> <li>• Write</li> </ul>                                                                                                         |

### Protocol Type NFS

| Field Name  | Description                                                                                                                                                                                                                                                                           |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Client Name | The name of the client that can access the NFS protocol. If multiple clients have been defined, an asterisk (*) is shown.                                                                                                                                                             |
| Read Only   | Specifies whether the export is read only. If the export is read only, files cannot be modified or created on this export. Values are: <ul style="list-style-type: none"> <li>• Yes (files cannot be modified or created)</li> <li>• No (files can be modified or created)</li> </ul> |

| Field Name  | Description                                                                                                                                                                                                                                                                                                     |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Root Squash | Specifies whether root squash is enabled for the export. <i>Root squash</i> refers to preventing root users on clients from having privileged access to exported files. Values are: <ul style="list-style-type: none"> <li>• Yes (root squash is enabled)</li> <li>• No (root squash is not enabled)</li> </ul> |
| Secure      | Specifies whether NFS access to the export is secure. Values are: <ul style="list-style-type: none"> <li>• Yes (NFS is secure)</li> <li>• No (NFS is not secure)</li> </ul>                                                                                                                                     |

*Storwize V7000U - Monitored Directory Information:*

This page shows general information about a directory that is associated with a Storwize V7000 Unified storage system.

The following information is displayed on this page:

| Field Name                                               | Description                                                                                                                                                                           |
|----------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Computer                                                 | The name of the computer that contains the directory.                                                                                                                                 |
| Directory                                                | The name of the directory.                                                                                                                                                            |
| Owner                                                    | The owner of the directory.                                                                                                                                                           |
| OS Group                                                 | The operating system group that the directory belongs to.                                                                                                                             |
| Total Size                                               | The total size of the directory.                                                                                                                                                      |
| File Count                                               | The number of files in the directory.                                                                                                                                                 |
| Directory Count                                          | The number of subdirectories in the directory.                                                                                                                                        |
| Average File Size                                        | The average size of the files in the directory.                                                                                                                                       |
| Access Time                                              | The date and time that the directory was last accessed.                                                                                                                               |
| Modification Time                                        | The date and time that the directory was last modified.                                                                                                                               |
| Creation Time                                            | The date and time that the directory was created.                                                                                                                                     |
| Scan Time                                                | The date and time that the directory was last scanned.                                                                                                                                |
| Discovered Time                                          | The date and time that the directory was discovered.                                                                                                                                  |
| Removed Time                                             | The date and time that the directory was removed.                                                                                                                                     |
| File attributes that are set                             | If any attributes were set for the files in the directory (for example, read only), the attributes are displayed. If no attributes were set, No File Attributes Are Set is displayed. |
| This Directory belongs to the following Directory Groups | A list of the directory groups that the directory belongs to.                                                                                                                         |

## Availability Reports

Use Availability reports to monitor and report on the availability of your storage, both from a network point of view and from a computer uptime perspective.

Use these reports to view the percentage of time that the storage in your enterprise is off-network due to network problems or is unavailable due to system downtime. In a fast-paced business environment, 24x7 application availability is critical. These reports help you monitor the availability of the storage and data that back up your applications.

---

| Reporting Type | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Ping           | <p>Use the ping report to monitor how available your storage is on the network. To do this monitor, regularly scheduled pings are sent out from the Data server to all of the monitored computers in your environment. The Tivoli Storage Productivity Center server checks not only for whether a Storage Resource agent or Data agent is available on the network, but records the amount of time that the request was in transit. This information indicates the responsiveness of the various storage servers on the network as well.</p> <p>Use the <b>Report Filter Specification</b> section on the <b>Selection</b> panel to limit the time range of the information in a report. Check the <b>Limit To This Time Range</b> option and use the date selection fields to limit the time range for a report.</p> <p>A ping report shows the following information:</p> <ul style="list-style-type: none"><li>• the name of a computer</li><li>• percentage of time the computer is available on the network</li><li>• number of pings that have been transmitted and the number that have been received</li><li>• number of damaged requests and the average transit time for the request to go between the Tivoli Storage Productivity Center server and the computer that is being monitored</li><li>• maximum transit time</li><li>• minimum transit time</li><li>• report start time and report end time for each computer in the report</li></ul> <p><b>Note:</b> You can sort on any of these columns by clicking the column name at the top of the report. You can also generate charts and drill down on each computer for more detail.</p> |

---

| Reporting Type  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Computer Uptime | <p data-bbox="737 222 1414 363">Use the computer uptime report to track the availability of the storage based upon the percentage of the time the computer that hosts that storage is down or not running. The Storage Resource agent or Data agent on the computer tracks this information and sends it to the server.</p> <p data-bbox="737 394 1386 420">A computer uptime report shows the following information:</p> <ul data-bbox="737 430 1398 699" style="list-style-type: none"> <li data-bbox="737 430 980 455">• name of a computer</li> <li data-bbox="737 466 1317 491">• percentage of time the computer is up and running</li> <li data-bbox="737 501 1227 527">• number of reboots that have been recorded</li> <li data-bbox="737 537 1357 594">• total amount of time the computer has been down and unavailable</li> <li data-bbox="737 604 1214 630">• average downtime that has been recorded</li> <li data-bbox="737 640 1398 699">• report start time and report end time for each computer in the report</li> </ul> <p data-bbox="737 720 1414 919"><b>Note:</b> Computer Uptime reports are not available for Storage Resource agents that are invoked using the non-daemon protocol for data collection. These reports do not contain data for computers on which non-daemon based Storage Resource agents are deployed. To collect computer uptime information using a Storage Resource agent, that agent must use a daemon service for runtime operation.</p> <p data-bbox="737 951 1414 1031">You can sort on any of these columns by clicking the column name at the top of the report. You can also generate charts and drill down on each computer for more detail.</p> |

### TPC-wide Storage Space Reports

Use these reports to view storage capacity of your resources at a file system, computer, or entire network level.

For a specific computer, group of computers, or all the computers in your enterprise, you can view the following:

- How much storage you have
- How much storage is being used
- What percentage of storage is free for expansion

| Reporting Type              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Disk Space                  | <p>Use these reports to view and chart disk capacity, per disk, per computer, per computer group, per domain, and for the whole network. These reports include one row for each computer that shows the total amount of storage capacity on that computer and the associated computer and disk storage information.</p> <p>For detail on any computer in the report, click the magnifying glass on the left of the report row. You can also click the chart button within a report to view an easy-to-interpret bar chart, showing the total storage capacity of each computer that you have chosen to report on and chart. To get the exact amount of storage for a computer, simply position the cursor over the associated bar and a box will appear showing the total amount for the selected computer.</p> <p><b>Note:</b> To ensure the accuracy and consistency of the data that appears in Disk Capacity reports, run regularly scheduled probes against the hosts that import volumes and the storage subsystems upon which you want to report. You should run probes against storage subsystems after running them against any hosts.</p> <p>You can use the following By clauses to view disk space for the following storage entities: By Disk, By Storage Subsystem, By Cluster, By Computer, By Computer Group, By Domain, and Network-wide.</p> |
| File System Space           | <p>Use these reports to view information for monitored file systems, including the total amount of storage capacity on those file systems and statistics for used space, free space, and total number of files. You can use the following By clauses to view file system space for the following storage entities: By Filesystem, By Filesystem Group, By Cluster, By Computer, By Computer Group, By Domain, Network-wide.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Consumed File System Space  | <p>Use these reports to view information about the amount of space consumed on a filesystem. You can use the following By clauses to view consumed file system space for the following storage entities: By Filesystem, By Filesystem Group, By Cluster, By Computer, By Computer Group, By Domain, Network-wide.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Available File System Space | <p>Use Filesystem Used Space reports to view information about the amount of free space on a filesystem. You can use the following By clauses to view available file system space for the following storage entities: By Filesystem, By Filesystem Group, By Cluster, By Computer, By Computer Group, By Domain, Network-wide.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

### TPC-Wide Storage Space - Disk Space report:

Use this report to view and chart disk storage capacity at the following levels: by disk, by computer, by cluster, by storage system, by computer monitoring or reporting group, by storage system monitoring or reporting group, by domain, and for the whole network.

**Note:** The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

#### By Disk

**Chart icon** 

Click this button to generate charts about the capacity of disks that are displayed in the report.

**Magnifying glass icon** 

Click this button to view a more detailed report about the capacity of a disk.

**Device**

Name of the device on which a disk resides. The following disks are displayed in this report:

- Computer disks that are internal to the computers monitored by IBM Tivoli Storage Productivity Center.
- Computer disks on virtual machines that are monitored by Tivoli Storage Productivity Center.
- Storage system disks on storage subsystems that are monitored by Tivoli Storage Productivity Center.
- Storage system disks that no longer exist after a probe of a storage system because information about that system was removed based on settings in the Removed Resource Retention panel.
- Back-end storage volumes that are not identified by Tivoli Storage Productivity Center as being assigned to a monitored storage system virtualizer.
- Computer disks that are not identified by Tivoli Storage Productivity Center as being assigned to a monitored storage system.

The following disks are not displayed in this report:

- Computer disks that are assigned to storage subsystems monitored by Tivoli Storage Productivity Center.
- Back-end storage volumes that are assigned to storage system virtualizers monitored by Tivoli Storage Productivity Center.

**Path** Name of a disk.

**Disk Space (Capacity)**

Total storage capacity for a disk. Keep in mind the following information about the **TOTAL** row for this column:

- The value that appears in **TOTAL** row for this column includes the disk space of spare disks within a storage subsystem.
- The value that appears in the **TOTAL** row for this column does not include computer disk space that is on virtual machines monitored by Tivoli Storage Productivity Center.

**Available Disk Space (Unallocated Space)**

Amount of space on the disk that is not allocated to any logical volume (or volumes for disks within storage subsystems). This value does not include RAID overhead on storage subsystems that are identified by Tivoli Storage Productivity Center. The value that appears in the **TOTAL** row for this column includes the unallocated capacity of all disks within the storage system, including the capacity of spare disks.

**In Storage Subsystem**

Indicates **Yes** when a disk is either of the following:

- An actual physical disk within a storage system that is monitored by Tivoli Storage Productivity Center.
- A back-end storage volume within a storage system virtualizer that is monitored by Tivoli Storage Productivity Center.

**Manufacturer**

Name of the manufacturer of a disk.

**Note:** The value **Unknown** is displayed in this column for IBM ESS storage subsystems.



**Model**

Model number of a disk.

**Note:** The value **Unknown** is displayed in this column for IBM ESS storage subsystems.

**Serial Number**

Serial number of a disk.

**Primary Defects**

Number of defects on the disk when it was new.

**Note:** Defect data is not collected for IBM TotalStorage Enterprise Storage Server, IBM FAStT, HP StorageWorks, Hitachi Data Systems storage subsystems.

**Grown Defects**

Number of defects detected on the disk since it was new.

**Note:** Defect data is not collected for TotalStorage Enterprise Storage Server, FAStT, HP StorageWorks, Hitachi Data Systems storage subsystems.

**Fibre Attached**

Indicates **Yes** when a disk is attached to a port that Tivoli Storage Productivity Center identifies as fibre.

**Mapped to Storage Subsystem**

Indicates **Yes** when a computer disk or a back-end storage volume can be mapped to a storage system.

**Detectable**

Indicates **Yes** when a disk is identified as being available to the server or the storage system following a successful probe.

**Is Virtual**

Indicates **Yes** when a disk resides on a computer that Tivoli Storage Productivity Center identifies as a virtual machine.

**Is Solid State**

Indicates **Yes** when a disk is on a solid-state drive.

**Is Encryptable**

Indicates **Yes** when a disk is encryptable.

**Is Encrypted**

Indicates **Yes** when a disk is encrypted.

**Tier** Indicates the tier level of a disk. Disks can have only one tier. For example, this column displays either the value **generic\_ssd** or **generic\_hdd**.

**By Storage Subsystem****Chart icon** 

Click this button to generate charts about the disk capacity of storage subsystems that are displayed in the report.

**Magnifying glass icon** 

Click this button to view a more detailed report about the capacity of disks associated with a storage system.

**Name**

Name or alias of a storage system that contains the disks or the name or

alias of a storage system virtualizer (for example, SAN Volume Controller) that manages the disks. This report does not display storage subsystems that have been discovered but not probed.

#### **Disk Space (Disk Capacity)**

Total raw (unformatted) disk capacity of a storage system. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Capacity of spare disks identified on TotalStorage Enterprise Storage Server, DS8000, and IBM System Storage DS6000 storage subsystems.
- Capacity of storage system disks that become missing after a Tivoli Storage Productivity Center storage system probe.
- The value that appears in the **TOTAL** row for this column does not include spare disks in its calculation.

#### **Available Disk Space (Disk Unallocated Space)**

Total unformatted disk freespace of a storage system. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this column:

- Storage system disks that become missing after a storage system probe.
- RAID overhead on storage subsystems that are identified by Tivoli Storage Productivity Center.
- External back-end storage for TagmaStore subsystems.
- Storage space that is consumed by volumes that reside on external disks when Hitachi Data Systems (HDS) storage subsystems are configured as storage virtualizers.
- The value that appears in the **TOTAL** row for this column does not include spare disks in its calculation.

#### **Number of Disks**

Total number of disks in a storage system. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Storage system disks that become missing after a storage system probe.
- External back-end storage for TagmaStore subsystems.

**Note:** If you click the magnifying glass icon next to storage system, the resulting drill-down report generates a list of detectable storage system disks or back-end storage volumes only. Therefore, the total number of disks for each storage system drill-down report will equal this value.

#### **Number of Volumes (LUNs)**

Total number of storage volumes in a storage system or the total number of virtual volumes that are virtual storage. This value does not include storage system volumes or virtual storage volumes that become missing after a storage system probe.

#### **Volume Space (LUN Capacity)**

Total capacity of all storage system volumes or virtual storage volumes within a storage system. This value does not include capacity information for storage system volumes or virtual storage volumes that become missing after a storage system probe.

**Note:** For space-efficient FlashCopy volumes, the volume capacity is the amount of storage space requested for these volumes, not the actual allocated amount. This can result in discrepancies in the overall storage

capacity reported for a storage system using space-efficient FlashCopy volumes. This also applies to the Total Usable Volume Space and Total FlashCopy Target Capacity calculations.

### **Consumable Volume Space (Usable LUN Capacity)**

Total amount of unique storage system volume space and virtual storage volume space on monitored storage system arrays that you can assign to servers within the network or are already assigned to servers within the network.

Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a probe.
- Storage system volumes or virtual storage volumes that are used as flash copy targets.
- Storage system volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
- Storage system volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database.
- Storage system volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage system virtualizers. For example, SAN Volume Controller.

**Note:** For FASTT, the XIV system, HP StorageWorks, and Hitachi Data Systems storage subsystems, the value in this column is the same as the value in the Volume Capacity column.

### **FlashCopy Target Volume Space (Capacity)**

Total amount of unique storage system volume space and virtual storage volume space on monitored storage system arrays that is identified as flash copy target space. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes or virtual storage volumes that are not used as flash copy targets.

### **Correlated Volume Space (LUN Capacity Visible to Data Manager)**

Total amount of storage system volume space or virtual storage volume space on systems where the systems disks are identified by Tivoli Storage Productivity Center as residing within a monitored storage system. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Disks that are assigned to a monitored storage system but are not identified by Tivoli Storage Productivity Center as a disk residing within the storage system. Some reasons why these disks are not identified as residing within the storage system include:
  - Tivoli Storage Productivity Center does not support the identification of disks for some storage system types.
  - The connection of the disk to the storage system is lost and is no longer detectable during a probe.

- Storage system volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
- Storage system volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database.
- Storage system volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage system virtualizers. For example, SAN Volume Controller.

**Volume Space Not Correlated (LUN Capacity not Visible to Data Manager)**

Total amount of unique storage system volume space and virtual storage volume space discovered on monitored storage system arrays that can be assigned or are assigned to systems within the network. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Disks that reside within a monitored storage system.
- Storage system volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage system virtualizers.

**Unavailable Disk Space (Overhead)**

Amount of storage system volume or virtual storage volume space that is dedicated to redundancy. This value is dependent on the storage system SMI-S provider returning a valid RAID value that is used to determine the overhead. If the value is zero for a storage system, the overhead cannot be calculated.

This amount does not include storage space information from storage system volumes or virtual storage volumes that become missing after a storage system probe.

**Formatted Space (Capacity)**

Total amount of formatted storage space associated with the storage system that is used or can be used for system storage volumes or virtual storage volumes. This value is dependent on the storage system SMI-S provider returning a formatted value.

For some storage subsystems, this is the total managed space which includes space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

For SAN Volume Controller, this value represents the total size of the managed disk groups created on the SAN Volume Controller. This value does not include managed disk space that is not defined to a managed disk group.

**Formatted Space with No Volumes**

Total amount of formatted storage space associated with the storage system that can be used for system storage volumes or virtual storage volumes. This value is dependent on the storage system SMI-S provider returning a formatted value.

For some subsystems, this value represents the remaining managed space that includes space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

For SAN Volume Controller, this value represents the space within the managed disk groups that does not include the size of the virtual disks created and the corresponding virtual disk extent overhead.

**Overall Unavailable Disk Space (Total Overhead)**

(TotalStorage Enterprise Storage Server, DS6000, and DS8000 storage systems only) Total amount of unformatted space within disk groups or array sites.

**Unformatted Disk Space**

This value does not apply to all storage subsystems. For FAS*t*T, TotalStorage Enterprise Storage Server, DS6000, and DS8000 storage systems, this is the total raw disk space of any unformatted disk groups or array sites. This column displays **N/A** for the storage systems where this value is not applicable. For the XIV system, this value shows the physical (hard) space of the XIV system available for pool creation or expansion.

**Cache**

Amount of cache memory that is internal to the storage system.

**Note:** This column is blank for SAN Volume Controllers.

**Manufacturer**

Name of the manufacturer of a disk.

**Model**

Model number of a disk.

**Serial Number**

Serial number of a disk. This information is displayed for physical disks only. For MDisks, the value **Unknown** is displayed.

**Firmware Revision**

Microcode (OS) version.

**Assigned Volume Space (Subsystem Assigned LUN Capacity)**

Total storage system volume space within the storage system that is mapped or assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes that are not mapped or are not assigned to host systems.

**Unassigned Volume Space (Subsystem Unassigned LUN Capacity)**

Total storage system volume space within a storage system that is not mapped or not assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes that are mapped or assigned to host systems.

**Unassigned and Available Formatted Disk Space (Subsystem Total Formatted Free Space)**

Total storage system space that is not mapped or not assigned to host systems and the total amount of formatted storage space associated with the storage system that can be used for system storage volumes.

This total amount of formatted storage space is dependent on the storage system SMI-S provider returning a formatted value. For some subsystems, this value represents the remaining managed space that might include space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes that become missing after a storage system probe.
- Storage system volumes that are mapped or assigned to host systems.

#### **zOS Volume Space (zOS LUN Capacity)**

Total storage system space that is reserved for use on mainframe storage such as used by z/OS. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes that become missing after a storage system probe.
- Storage system volumes that are reserved for use on open system server or non-mainframe storage.

#### **Open System Volume Space (Open System LUN Capacity)**

Total storage system space that is reserved for use on open system storage such as used by operating systems like LINUX, AIX, HP/UX, Solaris, and Windows. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes that become missing after a storage system probe.
- Storage system volumes that are reserved for use on mainframe storage.

#### **Physical Disk Space (Capacity)**

Total amount of physical disk space discovered on the monitored storage system arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system disks that become missing after a storage system probe.
- Back-end storage volumes on monitored storage system virtualizers (for example, SAN Volume Controller) that are attached from a storage system.

#### **Backend Volume Space (Backend LUN Capacity)**

Total amount of storage system volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on a monitored storage system virtualizer (for example, SAN Volume Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes or virtual storage volumes that are back-end disks on a storage system virtualizer that is not monitored.

#### **Configured Real Space (GB)**

Total amount of real available and formatted storage space in the storage pools that are associated with a storage subsystem.

**Available Real Space (GB)**

Amount of real storage space available to allocate to volumes in the storage pools that are associated with a storage subsystem.

**Tiers** Indicates the tiers that are used by the storage system. Each tier is delimited by a semi-colon (;). Possible values are: generic\_hdd, generic\_ssd

Information about tiers is available for SAN Volume Controller 6.1 and Storwize V7000 only. The value N/A is displayed in tier-related columns for other storage systems.

**Tier Capacities**

Total storage capacity for each tier that is used by the storage system. The capacity for each tier is delimited by a semi-colon (;).

**Tier Free Capacities**

Free storage capacity (available space) for each tier that is used by the storage system. The free capacity for each tier is delimited by a semi-colon (;).

**By Storage Subsystem Group****Chart icon** 

Click this button to generate charts about the disk capacity of storage subsystems in the groups that are displayed in the report.

**Magnifying glass icon** 

Click this button to view a more detailed report about the disk capacity of storage subsystems in a group.

**Creator**

ID of the user who created the storage system monitoring group or storage system reporting group.

**Group** Name of a group.

**Disk Space (Disk Capacity)**

Total raw (unformatted) disk capacity of the storage subsystems that are in a group. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Capacity of spare disks identified on TotalStorage Enterprise Storage Server, DS8000, and DS6000 storage systems.
- Capacity of storage system disks that become missing after a Tivoli Storage Productivity Center storage system probe.

**Available Disk Space (Disk Unallocated Space)**

Total unformatted disk free space of the storage subsystems that are in a group. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this column:

- Storage system disks that become missing after a storage system probe.
- RAID overhead on storage subsystems that are identified by Tivoli Storage Productivity Center.
- External back-end storage for TagmaStore subsystems.
- Storage space that is consumed by volumes that reside on external disks when Hitachi Data Systems (HDS) storage subsystems are configured as storage virtualizers.

**Number of Disks**

Total number of disks in the storage subsystems that are in a group. Tivoli

Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Storage system disks that become missing after a storage system probe.
- External back-end storage for TagmaStore subsystems.

**Note:** If you click the magnifying glass icon next to storage system, the resulting drill-down report generates a list of detectable storage system disks or back-end storage volumes only. Therefore, the total number of disks for each storage system drill-down report will equal this value.

#### **Number of Volumes (LUNs)**

Total number of storage volumes in a storage system or the total number of virtual volumes that are virtual storage. This value does not include storage system volumes or virtual storage volumes that become missing after a storage system probe.

#### **Volume Space (LUN Capacity)**

Total capacity of all storage system volumes or virtual storage volumes within the storage system. This value does not include capacity information of storage system volumes or virtual storage volumes that become missing after a storage system probe.

**Note:** For space-efficient FlashCopy volumes, the volume capacity is the amount of storage space requested for these volumes, not the actual allocated amount. This can result in discrepancies in the overall storage capacity reported for a storage system using space-efficient FlashCopy volumes. This also applies to the Total Usable Volume Space and Total FlashCopy Target Capacity calculations.

#### **Consumable Volume Space (Usable LUN Capacity)**

Total amount of unique storage system volume space and virtual storage volume space on monitored storage system arrays in a group that you can assign to servers within the network or are already assigned to servers within the network.

Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a probe.
- Storage system volumes or virtual storage volumes that are used as flash copy targets.
- Storage system volumes or virtual storage volumes that are identified as Business Continuity Volume extended (BCVx).
- Storage system volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database.
- Storage system volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage system virtualizers. For example, SAN Volume Controller.

**Note:** For FASTT, the XIV system, HP StorageWorks, and Hitachi Data Systems storage subsystems, the value in this column is the same as the value in the Volume Space column.

#### **FlashCopy Target Volume Space (Capacity)**

Total amount of unique storage system volume space and virtual storage volume space on monitored storage system arrays in a group that is



identified as flash copy target space. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes or virtual storage volumes that are not used as flash copy targets

#### **Correlated Volume Space (LUN Capacity Visible to Data Manager)**

Total amount of storage system volume space or virtual storage volume space on systems where the systems disks are identified by Tivoli Storage Productivity Center as residing within the monitored storage subsystems in a group. Tivoli Storage Productivity Center does not include the following scenarios in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Disks that are assigned to a monitored storage system but are not identified by Tivoli Storage Productivity Center as a disk residing within the storage system. Some reasons why these disks are not identified as residing within the storage system include:
  - Tivoli Storage Productivity Center does not support the identification of disks for some storage system types.
  - The connection of the disk to the storage system is lost and is no longer detectable during a probe.
- Storage system volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
- Storage system volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database.
- Storage system volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage system virtualizers. For example, SAN Volume Controller.

#### **Volume Space Not Correlated (LUN Capacity not Visible to Data Manager)**

Total amount of unique storage system volume space and virtual storage volume space discovered on monitored storage system arrays in a group that can be assigned or are assigned to systems within the network. Tivoli Storage Productivity Center does not include the following in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Disks that reside within a monitored storage system.
- Storage system volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage system virtualizers.

#### **Unavailable Disk Space**

Amount of volume or virtual storage volume space of storage subsystems in a group that is dedicated to redundancy. This value is dependent on the storage system SMI-S provider returning a valid RAID value that is used to determine the overhead. If the value is zero for a storage system, the overhead cannot be calculated.

This amount does not include storage space information from storage system volumes or virtual storage volumes that become missing after a storage system probe.

### **Formatted Space**

Total amount of formatted storage space associated with the storage subsystems in a group that is used or can be used for system storage volumes or virtual storage volumes. This value is dependent on the storage system SMI-S provider returning a formatted value.

For some storage subsystems, this is the total managed space which includes space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

For SAN Volume Controller, this value represents the total size of the managed disk groups created on the SAN Volume Controller. This value does not include managed disk space that is not defined to a managed disk group.

### **Formatted Space with No Volumes**

Total amount of formatted storage space associated with the storage subsystems in a group that can be used for system storage volumes or virtual storage volumes. This value is dependent on the storage system SMI-S provider returning a formatted value.

For some subsystems, this value represents the remaining managed space that includes space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

For SAN Volume Controller, this value represents the space within the managed disk groups that does not include the size of the virtual disks created and the corresponding virtual disk extent overhead.

### **Unformatted Disk Space**

This value does not apply to all storage subsystems. For FAS*t*T, TotalStorage Enterprise Storage Server, DS6000, and DS8000 storage systems, this is the total raw disk space of any unformatted disk groups or array sites. For the XIV system, the total available physical (hard) space that is left for pool creation. This column displays *N/A* for the storage subsystems where this value is not applicable.

### **Cache**

Amount of cache memory that is internal to the storage subsystems in the group.

**Note:** This column does not include values for SAN Volume Controllers.

### **Assigned Volume Space (Subsystem Assigned LUN Capacity)**

Total volume space of the storage subsystems within a group that is mapped or assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes that are not mapped or are not assigned to host systems.

### **Unassigned Volume Space (Subsystem Unassigned LUN Capacity)**

Total volume space of the storage subsystems in a group that is not mapped or not assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes that are mapped or assigned to host systems.

#### **zOS Volume Space (zOS LUN Capacity)**

Total space of storage subsystems in a group that is reserved for use on mainframe storage such as used by z/OS. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes that become missing after a storage system probe.
- Storage system volumes that are reserved for use on open system server or non-mainframe storage.

#### **Open System Volume Space (Open System LUN Capacity)**

Total space of the storage subsystems in a group that is reserved for use on open system storage such as used by operating systems such as Linux, AIX, HP/UX, Solaris, and Windows. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes that become missing after a storage system probe.
- Storage system volumes that are reserved for use on mainframe storage.

#### **Physical Disk Space (Capacity)**

Total amount of physical disk space on the storage system arrays that are in a group. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system disks that become missing after a storage system probe.
- Back-end storage volumes on monitored storage system virtualizers (for example, SAN Volume Controller) that are attached from a storage system.

#### **Backend Volume Space (Backend LUN Capacity)**

Total amount of storage system volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on the monitored storage system virtualizers (for example, SAN Volume Controller) that are in a group. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes or virtual storage volumes that are back-end disks on a storage system virtualizer that is not monitored.

#### **By Cluster**

##### **Chart icon**

Click this button to generate charts about the disk capacity of the clusters that are displayed in the report.

##### **Magnifying glass icon**

Click this button to view a more detailed report about the capacity of disks in a cluster.

**Cluster**

Name of a cluster.

**Disk Space (Total Capacity)**

Total amount of unique disk storage space for the computers in a cluster. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value in this column:

- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

**Note:**

- If the computer is an IBM SAN File System client, then this value will also include the capacity of any storage system volumes visible to this computers that are owned by the SAN File System.
- The drill-down reports for each computer generates a list of detectable computer disks only. Therefore, the total capacity of disks for each computer drill-down report will equal this value.

**Available Disk Space (Total Unallocated Space)**

Total amount of unique disk storage space for a computer that is not allocated to any logical volume on the computer. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

**Note:** If the computer is a SAN File System client, then this value does not include the unallocated space on any storage system volumes visible to it that are owned by the SAN File System.

**Owned Disk Space (Capacity)**

Total amount of unique disk storage space for computers in a cluster that is owned by those computers. If Data Manager determines that a disk is configured for a different host, that disk's capacity is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

**Note:**

- If the computer is a SAN File System client, this value does not include the capacity of any volumes that are visible to it that are owned by the SAN File System.
- If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and reads from the host ID. If this does not agree with the local host ID, the disk is classified as belonging to a different (Solaris) host.

### **Owned Available Disk Space (Owned Unallocated Space)**

Total amount of unique disk storage space for computers in a cluster that is not allocated to any logical volume within the computers and is owned by the computers. If Data Manager determines that a disk is configured for a different computer, that disk's unallocated space is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

If the computer is a SAN File System client, then this value excludes the unallocated space on any storage system volumes visible to it that are owned by the SAN File System.

If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product will examine the Veritas disk-label and read off the host-ID. If this does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.

### **Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)**

Total amount of unique disk space on the computers in a cluster that is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that reside behind a fibre channel port.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

### **Non-Fibre Channel Attached Available Disk Space (Total Other Attached Unallocated Space)**

Total amount of disk storage space on the computers in a cluster that is not allocated to any logical volume within the computers and is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that reside behind a fibre channel port.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

### **Fibre Channel Attached Disk Space (Total Fibre Channel Attached Capacity)**

Total amount of unique disk space on the computers in a cluster that is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that do not reside behind a fibre channel port.
- Computer disks that reside behind a fibre channel port that is not identified by a Storage Resource agent or Data agent.
- Computer disks on a computer that has not been probed.

- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

#### **Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Unallocated Space)**

Total amount of disk storage space on the computers in a cluster that is not allocated to any logical volume within the computers and is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that do not reside behind a fibre channel port.
- Computer disks that reside behind a fibre channel port that is not identified by a Storage Resource agent or Data agent.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

#### **By Computer**

##### **Chart icon**

Click this button to generate charts about the disk capacity of the computers that are displayed in the report.

##### **Magnifying glass icon**

Click this button to view a more detailed report about the disks on a computer.

#### **Computer**

Name of a computer.

#### **Disk Space (Total Capacity)**

Total amount of unique disk storage space for a computer. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

#### **Note:**

- The value that appears in the **TOTAL** row for this column does not include disks on virtual machines where the Storage Resource agent or Data agent is installed.
- If the computer is a SAN File System client, then this value will also include the capacity of any storage system volumes visible to this computers that are owned by the SAN File System.
- The drill-down reports for each computer generates a list of detectable computer disks only. Therefore, the total capacity of disks for each computer drill-down report will equal this value.

#### **Available Disk Space (Total Unallocated Space)**

Total amount of unique disk storage space for a computer that is not allocated to any logical volume on the computer. Tivoli Storage

Productivity Center does not include the following disks from in its calculation of the value in this column:

- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

**Note:** If the computer is a SAN File System client, then this value does not include the unallocated space on any storage system volumes visible to it that are owned by the SAN File System.

#### **Owned Disk Space (Capacity)**

Total amount of unique disk storage space for a computer that is owned by that computer. If Data Manager determines that a disk is configured for a different host, that disk's capacity is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

If the computer is a SAN File System client, this value does not include the capacity of any volumes that are visible to it that are owned by the SAN File System.

If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and reads from the host ID. If this does not agree with the local host ID, the disk is classified as belonging to a different (Solaris) host.

#### **Owned Available Disk Space (Owned Unallocated Space)**

Total amount of unique disk storage space for a computer that is not allocated to any logical volume within the computer and is owned by the computer. If Data Manager determines that a disk is configured for a different computer, that disk's unallocated space is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

#### **Note:**

- If the computer is a SAN File System client, then this value excludes the unallocated space on any storage system volumes visible to it that are owned by the SAN File System.
- If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product will examine the Veritas disk-label and read off the host-ID. If this does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.

#### **Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)**

Total amount of unique disk space on a computer that is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that reside behind a fibre channel port.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

**Non-Fibre Channel Attached Available Disk Space (Total Other Attached Unallocated Space)**

Total amount of disk storage space for a computer that is not allocated to any logical volume within the computer and is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that reside behind a fibre channel port.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

**Fibre Channel Attached Disk Space (Total Fibre Channel Attached Capacity)**

Total amount of unique disk space on a computer that is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that do not reside behind a fibre channel port.
- Computer disks that reside behind a fibre channel port that is not identified by a Storage Resource agent or Data agent.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

**Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Unallocated Space)**

Total amount of disk storage space on a computer that is not allocated to any logical volume within the computer and is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that do not reside behind a fibre channel port.
- Computer disks that reside behind a fibre channel port that is not identified by a Storage Resource agent or Data agent.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

**OS Type**

Operating system running on a computer.

**Network Address**

Network address of a computer.

**IP Address**

IP address of a computer.

**Time Zone**

Time zone in which a computer is running.



**CPU Architecture**

Architecture of the processor (for example, IA32).

**By Computer Group****Chart icon** 

Click this button to generate charts about the disk capacity of computers in the groups that are displayed in the report.

**Magnifying glass icon** 

Click this button to view a more detailed report about the disk capacity of computers in a group.

**Creator**

Creator of a computer monitoring group or computer reporting group.

**Group**

Name of a group.

**Disk Space (Total Capacity)**

Total amount of unique disk storage space for the computers in a group. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that are on a virtual machine monitored by Tivoli Storage Productivity Center.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

**Available Disk Space (Total Unallocated Space)**

Total amount of unique disk storage space for computers in a group that is not allocated to any logical volume on the computers. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

**Note:** If the computer is a SAN File System client, then this value does not include the unallocated space on any storage system volumes visible to it that are owned by the SAN File System.

**Owned Disk Space (Capacity)**

Total amount of unique disk storage space for computers in a group that is owned by those computers. If Data Manager determines that a disk is configured for a different host, that disk's capacity is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on a cluster of virtual machines that is monitored by Tivoli Storage Productivity Center.
- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

**Note:**

- If the computer is a SAN File System client, this value does not include the capacity of any volumes that are visible to it that are owned by the SAN File System.
- If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and reads from the host ID. If this does not agree with the local host ID, the disk is classified as belonging to a different (Solaris) host.

**Owned Available Disk Space (Owned Unallocated Space)**

Total amount of unique disk storage space for the computers in a group that is not allocated to any logical volume within the computer group and is owned by a computer within the group. If Data Manager determines that a disk is configured for a different computer, that disk's unallocated space is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

**Note:**

- If the computer is a SAN File System client, then this value excludes the unallocated space on any storage system volumes visible to it that are owned by the SAN File System.
- If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product will examine the Veritas disk-label and read off the host-ID. If this does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.

**Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)**

Total amount of unique disk space on the computers in a group and is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that reside behind a fibre channel port.
- Computer disks discovered on virtual machine servers where the Storage Resource agent or Data agent is installed.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

**Non-Fibre Channel Attached Available Disk Space (Total Other Attached Unallocated Space)**

Total amount of disk storage space on the computers in a group that is not allocated to any logical volume within the computers and is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that reside behind a fibre channel port.
- Computer disks on a computer that has not been probed.

- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

#### **Fibre Channel Attached Disk Space (Total Fibre Channel Attached Capacity)**

Total amount of unique disk space on the computers in a group that is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that do not reside behind a fibre channel port.
- Computer disks that reside behind a fibre channel port that is not identified by a Storage Resource agent or Data agent.
- Computer disks discovered on virtual machine servers where the Storage Resource agent or Data agent is installed.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

#### **Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Unallocated Space)**

Total amount of disk storage space on the computers in a group that is not allocated to any logical volume within the computers and is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks that do not reside behind a fibre channel port.
- Computer disks that reside behind a fibre channel port that is not identified by a Storage Resource agent or Data agent.
- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

#### **By Domain**

##### **Chart icon**

Click this button to generate charts about the disk capacity of the computers in a domain.

##### **Magnifying glass icon**

Click this button to view a more detailed report about the disk capacity of computers in a domain.

#### **Domain**

Name of a domain.

#### **Disk Space (Capacity)**

Total amount of unique disk storage space for the computers in a domain. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

### **Available Disk Space (Unallocated Space)**

Total amount of unique disk storage space for computers in a domain that is not allocated to any logical volume on the computers. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on computers that have not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.

### **Network-wide**

#### **Chart icon**

Click this button to generate charts about the disk capacity across a network.

#### **Magnifying glass icon**

Click this button to view a more detailed report about the disk capacity of computers in a network.

### **Disk Space (Capacity)**

Total amount of internal physical disk space on monitored storage system arrays and disk space on monitored systems that are not storage system volumes. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage system disks or computer disks that become missing after a probe.
- Back-end storage volumes that are not identified by Tivoli Storage Productivity Center as being assigned to a monitored storage system virtualizer.
- Computer disks that are identified by Tivoli Storage Productivity Center as being assigned to a monitored storage system.
- Computer disks that reside within a virtual machine server where the Storage Resource agent or Data agent is installed.

### **Consumed File System Space (Filesystem Used Space)**

Total amount of unique file system storage space that is used or consumed by the systems where the Storage Resource agent or Data agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems on discovered hypervisors that are assigned to discovered virtual machines where the Storage Resource agent or Data agent is installed.
- File systems on discovered virtual machines where the Storage Resource agent or Data agent is installed, but the hypervisor for the virtual machine has not been discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

### **Available File System Space (Filesystem Free Space)**

Total amount of unique file system storage space that is not used or

available to the systems where the Storage Resource agent or Data agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

#### **Raw Volume Space**

Total amount of storage system volume space that is the following:

- Assigned or mapped to hosts.
- Storage system volume space on a system that it is monitored by Tivoli Storage Productivity Center.
- Logical volume space on hosts that do not reside within a storage system monitored by Tivoli Storage Productivity Center.

#### **Unavailable Disk Space (Overhead)**

Total amount of storage system volume space that is identified as the following:

- For RAID or mirror redundancy on a storage system.
- As storage system volume space used for replication targets.
- As logical volume space on hosts that is for RAID or mirror redundancy on those hosts.

#### **Available Disk Space (Unallocated Space)**

Total amount of disk space that does not include the following from the network capacity:

- Network Raw Volume Space
- Network Overhead
- Network Unknown Volume Space

#### **Unknown Storage Subsystem Volume Space (Unknown LUN Capacity)**

Total amount of storage system volume space that is not mapped or assigned to a host and is not identified as residing on a host that has been probed by a Storage Resource agent or Data agent.

#### **TPC-Wide Storage Space - File System Space report:**

Use these reports to view detailed information about the storage capacity of file systems on machines in your environment. You can view this information in different ways using the following By clauses: by filesystem, by filesystem group, by cluster, by computer, by computer group, by domain, and for the entire network.

**Note:** The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

#### **By Filesystem**

**Chart icon** 

Click this button to generate charts about the storage space for a file system.

**Computer**

Name of the computer on which a file system resides.

**Mount Point**

Name or mount point (UNIX/Linux) of the file system (for example, c:\, d;\, /opt, /export/home, etc.). The file systems displayed in this report include the following:

- File systems on virtual machines that are monitored by IBM Tivoli Storage Productivity Center.
- File systems on hypervisors that are monitored by Tivoli Storage Productivity Center.
- File systems that become missing after a probe.

**File System Space (Capacity)**

Amount of file system storage space for a file system.

**Note:**

- The value that appears in the **TOTAL** row for this column does not include file systems on virtual machines where the Storage Resource agent or Data agent is installed.
- The value in this column might be greater than **Used Space** + **Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the `df` command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space** + **Free Space**, the difference representing the space wasted due to the system overhead.

**Detectable**

Indicates **Yes** when a file system is identified as being available to the server following a successful probe.

**Percent Used Space**

Percentage of file system storage space that is used within a file system.

**Consumed File System Space (Used Space)**

Amount of used storage space within a file system. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead). Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

The value that appears in the **TOTAL** row for this column represents the difference between the **TOTAL** values in the **File System Space** and **Available File System Space** columns. If the calculated value for the **TOTAL** row is negative, the column displays 0.

The **TOTAL** row does not add the used space coming from VMWare computers. See the *VMWare capacity reports* topic in the *Planning* chapter of the *Tivoli Storage Productivity Center Installation and Configuration Guide* for information about how file system and disk capacity are calculated in this

report for VMWare computers. This topic is also available in the *Planning* section of the *Tivoli Storage Productivity Center Information Center*.

The value that appears in the **TOTAL** row for this column does not include the following file systems:

- File systems on discovered hypervisors that are assigned to virtual machine servers where the Storage Resource agent or Data agent has been installed.
- File systems on discovered virtual machine servers where the Storage Resource agent or Data agent is installed, but the hypervisor for the virtual machine servers has not been discovered and the total used space exceeds the capacity of the total file system capacity.

#### **Available File System Space (Free Space)**

Amount of unused storage space within a file system.

#### **File Count**

Total number of files within a file system. This value does not include files on file systems that have not been scanned.

**Note:** On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

#### **Directory Count**

Total number of directories within a file system. This value does not include directories on file systems that have not been scanned.

#### **Percent Used Inodes**

Percentage of used Inodes within a file system.

#### **Used Inodes**

Number of used Inodes within a file system.

#### **Free Inodes**

Number of free Inodes within a file system.

#### **Volume Group**

Volume group to which a logical volume belongs. If the volume is not part of a volume group this column displays **N/A**.

#### **Logical Volume Space (Logical Volume Allocation)**

Logical volume space used by the logical volume containing the file system.

#### **Disk Detectable**

Indicates **Yes** when a file system's logical volume is identified as having an underlying host disk.

#### **By Filesystem Group**

#### **Chart icon**

Click this button to generate charts about the capacity of file systems in file system monitoring groups or file system reporting groups.

### Magnifying glass icon

Click this button to view more detailed information about the capacity of file systems in a file system monitoring group or file system reporting group.

### Creator

ID of the user who created a group.

**Group** Name of a file system monitoring group or file system reporting group.

### File System Space (Capacity)

Amount of file system storage space for the file systems in a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on virtual machines where the Storage Resource agent or Data agent is installed.
- File systems that become missing after a probe.
- File systems that are not mounted.

**Note:** The value in this column might be greater than **Used Space + Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the `df` command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space + Free Space**, the difference representing the space wasted due to the system overhead.

### Percent Used Space

Percentage of total file system storage space that is used within a group.

### Consumed File System Space (Used Space)

Total amount of used file system storage space within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on hypervisors that are assigned to virtual machine servers where the Storage Resource agent or Data agent is installed if both the virtual machine file systems and hypervisor file systems reside within the group.
- File systems on virtual machine servers where the Storage Resource agent or Data agents are deployed, but the hypervisor for the virtual machine servers has not been discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems on virtual machine servers that are in a group without the underlying hypervisor file systems, and the total used space exceeds the capacity of the total file system capacity.
- File systems that become missing after a probe.
- File systems that are not mounted.

**Note:** On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

### Available File System Space (Free Space)

Total amount of unique file system storage space that is not used or is available within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:



- File systems that become missing after a probe.
- File systems that are not mounted.

#### **File Count**

Total number of files within a group. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Note:** On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

#### **Directory Count**

Total number of directories within a group. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Storage Resource agent or Data agent.

#### **Percent Used Inodes**

Percentage of used Inodes within a group.

#### **Used Inodes**

Number of used Inodes within a group.

#### **Free Inodes**

Number of free Inodes within a group.

#### **By Cluster**

##### **Chart icon**

Click this button to generate charts about the used file system space for a cluster.

##### **Magnifying glass icon**

Click this button to view more detailed information about the used storage space on file systems that are in a cluster.

#### **Cluster**

Name of cluster.

#### **File System Space (Capacity)**

Total amount of unique file system storage space in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.

- File systems that are not mounted.

**Note:** The value in this column might be greater than **Used Space + Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the `df` command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space + Free Space**, the difference representing the space wasted due to the system overhead.

#### **Percent Used Space**

Percentage of total file system storage space that is used within a cluster.

#### **Consumed File System Space (Used Space)**

Total amount of unique file system storage space that is used or consumed by the file systems in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**Note:** On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

#### **Available File System Space (Free Space)**

Total amount of unique file system storage space that is not used or is available in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

#### **File Count**

Total number of files within a cluster. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by an agent.

When monitoring computers where the Windows operating system is installed, the number of files indicated in this column might be lower than the number of files reported by the operating system.

#### **Directory Count**

Total number of directories within a cluster. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by an agent.

When monitoring computers where the Windows operating system is installed, the number of folders indicated in this column might be higher than the number of folders that is reported by the operating system.

#### **Percent Used Inodes**

Percentage of used Inodes within a cluster.

#### **Used Inodes**

Number of used Inodes within a cluster.

#### **Free Inodes**

Number of free Inodes within a cluster.

#### **By Computer**

#### **Chart icon**

Click this button to generate a chart about the used file system storage space of the computers that are displayed in the report.

#### **Magnifying glass icon**

Click this button to view a more detailed report about the used file system storage space on a computer.

#### **File System Space (Capacity)**

Total amount of unique file system storage space that was discovered on the computer where the Storage Resource agent or Data agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

#### **Note:**

- The value in this column might be greater than **Used Space + Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the `df` command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space + Free Space**, the difference representing the space wasted due to the system overhead.
- The value that appears in the **TOTAL** row for this column does not include file systems on virtual machines where the Storage Resource agent or Data agent is installed.
- The drill-down reports for each computer will generate a list of detectable file systems only. Therefore, the total capacity of file systems for each computer drill-down report will equal this value.

#### **Percent Used Space**

Percentage of total file system storage space that is used within a computer.

#### **Consumed File System Space (Used Space)**

Amount of unique file system storage space that is used or consumed by the computer where the Storage Resource agent or Data agent is installed. On some UNIX file systems, this value does not include storage space

reserved by the operating system (overhead). Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

The value that appears in the **TOTAL** row for this column represents the difference between the **TOTAL** values in the **File System Space** and **Available File System Space** columns. If the calculated value for the **TOTAL** row is negative, the column displays 0.

The **TOTAL** row does not add the used space coming from VMWare computers. See the *VMWare capacity reports* topic in the *Planning* chapter of the *Tivoli Storage Productivity Center Installation and Configuration Guide* for information about how file system and disk capacity are calculated in this report for VMWare computers. This topic is also available in the *Planning* section of the *Tivoli Storage Productivity Center Information Center*.

The value that appears in the **TOTAL** row for this column does not include the following file systems:

- File systems on discovered hypervisors that are assigned to virtual machine servers where the Storage Resource agent or Data agent has been installed.
- File systems on discovered virtual machine servers where the Storage Resource agent or Data agent is installed, but the hypervisor for the virtual machine servers has not been discovered and the total used space exceeds the capacity of the total file system capacity.

#### **Available File System Space (Free Space)**

Total amount of unique file system storage space that is not used or is available to a computer. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

#### **File Count**

Total number of files within a computer. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Note:** On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the **File Count** column is 200.

**Directory Count**

Total number of directories within a computer. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Percent Used Inodes**

Percentage of used Inodes within a computer.

**Used Inodes**

Number of used Inodes within a computer.

**Free Inodes**

Number of free Inodes within a computer.

**By Computer Group****Chart icon** 

Click this button to generate charts about the used file system storage space for a computer monitoring group or computer reporting group.

**Magnifying glass icon** 

Click this button to view more detailed information about the file systems in a computer monitoring group or computer reporting group.

**Creator**

ID of the user who created a group.

**Group** Name of a computer monitoring group or computer reporting group.

**File System Space (Capacity)**

Amount of file system storage space for the computers in a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on virtual machines where the Storage Resource agent or Data agent is installed.
- File systems that become missing after a probe.
- File systems that are not mounted.

**Note:** The value in this column might be greater than **Used Space + Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the `df` command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space + Free Space**, the difference representing the space wasted due to the system overhead.

**Percent Used Space**

Percentage of total file system storage space that is used within a group.

**Consumed File System Space (Used Space)**

Total amount of used file system storage space within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on hypervisors that are assigned to virtual machine servers where the Storage Resource agent or Data agent is installed if both the virtual machine file systems and hypervisor file systems reside within the group.
- File systems on virtual machine servers where the Storage Resource agent or Data agent is installed, but the hypervisor for the virtual machine servers has not been discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems on virtual machine servers that are in a group without the underlying hypervisor file systems, and the total used space exceeds the capacity of the total file system capacity.
- File systems that become missing after a probe.
- File systems that are not mounted.

**Note:** On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

#### **Available File System Space (Free Space)**

Total amount of unique file system storage space that is not used or is available within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are not mounted.

#### **File Count**

Total number of files within a group. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Note:** For groups that contain Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the **File Count** column is 200.

#### **Directory Count**

Total number of directories within a group. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Storage Resource agent or Data agent.

#### **Percent Used Inodes**

Percentage of used Inodes within a group.

#### **Used Inodes**

Number of used Inodes within a group.

**Free Inodes**

Number of free Inodes within a group.

**By Domain****Chart icon** 

Click this button to generate charts about the file system used space for a domain.

**Magnifying glass icon** 

Click this button to view more detailed information about the file system used space of computers that are in a domain.

**Domain**

Name of a domain.

**File System Space (Capacity)**

Total amount of unique file system storage space that was discovered on the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**Note:** The value in this column might be greater than **Used Space + Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the `df` command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space + Free Space**, the difference representing the space wasted due to the system overhead.

**Percent Used Space**

Percentage of total file system storage space that is used within a domain.

**Consumed File System Space (Used Space)**

Total amount of unique file system storage space that is used or consumed by the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**Note:** On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

**Available File System Space (Free Space)**

Total amount of unique file system storage space that is not used or is available to the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.

- File systems that are mounted remotely and are not identified to a network attached storage server
- File systems that are not mounted.

#### **File Count**

Total number of files within a domain. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Note:** On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

#### **Directory Count**

Total number of directories within a domain. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Storage Resource agent or Data agent.

#### **Percent Used Inodes**

Percentage of used Inodes within a domain.

#### **Used Inodes**

Number of used Inodes within a domain.

#### **Free Inodes**

Number of free Inodes within a domain.

#### **Network-wide**

##### **Chart icon**

Click this button to generate a chart about the file system used space in a network.

##### **Magnifying glass icon**

Click this button to view more detailed information about the file system used space in a network.

#### **File System Space (Capacity)**

Total amount of unique file system storage space that was discovered on the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.



- File systems that are not mounted.

**Note:** The value in this column might be greater than **Used Space + Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agent or Data agents and by the `df` command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space + Free Space**, the difference representing the space wasted due to the system overhead.

#### **Percent Used Space**

Percentage of total file system storage space that is used within a network.

#### **Consumed File System Space (Used Space)**

Total amount of unique file system storage space that is used or consumed by the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**Note:** On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

#### **Available File System Space (Free Space)**

Total amount of unique file system storage space that is not used or is available to the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server
- File systems that are not mounted.

#### **File Count**

Total number of files within a network. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Note:** On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

#### **Directory Count**

Total number of directories within a network. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Percent Used Inodes**

Percentage of used Inodes within a network.

**Used Inodes**

Number of used Inodes within a network.

**Free Inodes**

Number of free Inodes within a network.

**TPC-Wide Storage Space - Consumed File System Space report:**

Use these reports to view detailed information about the used storage space in file systems on computers in your environment. You can view this information in different ways using the following By clauses: by filesystem, by filesystem group, by cluster, by computer, by computer group, by domain, and for the entire network.

**Note:** The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

**By Filesystem**

**Chart icon** 

Click this button to generate charts about the used space information for a file system.

**Computer**

Name of the computer on which a file system resides.

**Mount Point**

Name or mount point (UNIX/Linux) of the file system (for example, c:\, d;\, /opt, /export/home, etc.). The file systems displayed in this report include the following:

- File systems on virtual machines that are monitored by IBM Tivoli Storage Productivity Center.
- File systems on hypervisors that are monitored by Tivoli Storage Productivity Center.
- File systems that become missing after a probe.

**Consumed File System Space (Used Space)**

Amount of used storage space within a file system. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

The value that appears in the **TOTAL** row for this column does not include the following file systems:

- File systems on discovered hypervisors that are assigned to virtual machine servers where the Storage Resource agent or Data agent has been installed.
- File systems on discovered virtual machine servers where the Storage Resource agent or Data agent is installed, but the hypervisor for the

virtual machine servers has not been discovered and the total used space exceeds the capacity of the total file system capacity.

**Detectable**

Indicates **Yes** when a file system is identified as being available to the server following a successful probe.

**Percent Used Space**

Percentage of file system storage space that is used within a file system.

**Available File System Space (Free Space)**

Amount of unused storage space within a file system.

**File System Space (Capacity)**

Amount of file system storage space for a file system.

**Note:**

- The value that appears in the **TOTAL** row for this column does not include file systems on virtual machines where the Storage Resource agent or Data agent is installed.
- The value in this column might be greater than **Used Space + Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the df command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space + Free Space**, the difference representing the space wasted due to the system overhead.

**File Count**

Total number of files within a file system. This value does not include files on file systems that have not been scanned.

**Note:** On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

**Directory Count**

Total number of directories within a file system. This value does not include directories on file systems that have not been scanned.

**Percent Used Inodes**

Percentage of used Inodes within a file system.

**Used Inodes**

Number of used Inodes within a file system.

**Free Inodes**

Number of free Inodes within a file system.

**Volume Group**

Volume group to which a logical volume belongs. If the volume is not part of a volume group this column displays **N/A**.

**Logical Volume Space (Allocation)**

Logical volume space used by the logical volume containing the file system.

**Disk Detectable**

Indicates **Yes** when a file system's logical volume is identified as having an underlying host disk.

**By Filesystem Group****Chart icon** 

Click this button to generate charts about the used file system storage space in a file system monitoring group or file system reporting group.

**Magnifying glass icon** 

Click this button to view more detailed information about the used file system storage space in a file system monitoring group or file system reporting group.

**Creator**

ID of the user who created a group.

**Group** Name of a file system monitoring group or file system reporting group.

**Consumed File System Space (Used Space)**

Total amount of used file system storage space within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on hypervisors that are assigned to virtual machine servers where the Storage Resource agent or Data agent is installed if both the virtual machine file systems and hypervisor file systems reside within the group.
- File systems on virtual machine servers where the Storage Resource agent or Data agent is installed, but the hypervisor for the virtual machine servers has not been discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems on virtual machine servers that are in a group without the underlying hypervisor file systems, and the total used space exceeds the capacity of the total file system capacity.
- File systems that become missing after a probe.
- File systems that are not mounted.

**Note:** On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

**Percent Used Space**

Percentage of total file system storage space that is used within a group.

**Available File System Space (Free Space)**

Total amount of unique file system storage space that is not used or is available within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are not mounted.

**File System Space (Capacity)**

Amount of file system storage space for the file systems in a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on virtual machines where the Storage Resource agent or Data agent is installed.
- File systems that become missing after a probe.
- File systems that are not mounted.

**Note:** The value in this column might be greater than **Used Space + Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the `df` command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space + Free Space**, the difference representing the space wasted due to the system overhead.

#### **File Count**

Total number of files within a group. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Note:** On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

#### **Directory Count**

Total number of directories within a group. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Storage Resource agent or Data agent.

#### **Percent Used Inodes**

Percentage of used Inodes within a group.

#### **Used Inodes**

Number of used Inodes within a group.

#### **Free Inodes**

Number of free Inodes within a group.

#### **By Cluster**

#### **Chart icon**

Click this button to generate charts about the used file system space for a cluster.

#### **Magnifying glass icon**

Click this button to view more detailed information about the used storage space on file systems that are in a cluster.

**Cluster**

Name of cluster.

**Consumed File System Space (Used Space)**

Total amount of unique file system storage space that is used or consumed by the file systems in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**Note:** On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

**Percent Used Space**

Percentage of total file system storage space that is used within a cluster.

**Available File System Space (Free Space)**

Total amount of unique file system storage space that is not used or is available in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**File System Space (Capacity)**

Total amount of unique file system storage space in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**Note:** The value in this column might be greater than **Used Space + Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the `df` command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space + Free Space**, the difference representing the space wasted due to the system overhead.

**File Count**

Total number of files within a cluster. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Storage Resource agent or Data agent.

When monitoring computers where the Windows operating system is installed, the number of files indicated in this column might be lower than the number of files reported by the operating system.

#### **Directory Count**

Total number of directories within a cluster. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Storage Resource agent or Data agent.

When monitoring computers where the Windows operating system is installed, the number of folders indicated in this column might be higher than the number of folders that is reported by the operating system.

#### **Percent Used Inodes**

Percentage of used Inodes within a cluster.

#### **Used Inodes**

Number of used Inodes within a cluster.

#### **Free Inodes**

Number of free Inodes within a cluster.

#### **By Computer**

##### **Chart icon**

Click this button to generate charts about the used file system storage space of the computers that are displayed in the report.

##### **Magnifying glass icon**

Click this button to view a more detailed report about the used file system storage space on a computer.

#### **Computer**

Name of the computer on which a file system resides.

#### **Consumed File System Space (Used Space)**

Total amount of unique file system storage space that is used or consumed by the computer where the Storage Resource agent or Data agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**Note:** On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

#### **Percent Used Space**

Percentage of total file system storage space that is used within a computer.

#### **Available File System Space (Free Space)**

Total amount of unique file system storage space that is not used or is

available to a computer. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

#### **File System Space (Capacity)**

Total amount of unique file system storage space that was discovered on the computer where the Storage Resource agent or Data agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

#### **Note:**

- The value in this column might be greater than **Used Space + Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the `df` command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space + Free Space**, the difference representing the space wasted due to the system overhead.
- The value that appears in the **TOTAL** row for this column does not include file systems on virtual machines where the Storage Resource agent or Data agent is installed.
- The drill-down reports for each computer will generate a list of detectable file systems only. Therefore, the total capacity of file systems for each computer drill-down report will equal this value.

#### **File Count**

Total number of files within a computer. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Note:** On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

#### **Directory Count**

Total number of directories within a computer. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:



- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Percent Used Inodes**

Percentage of used Inodes within a computer.

**Used Inodes**

Number of used Inodes within a computer.

**Free Inodes**

Number of free Inodes within a computer.

**By Computer Group**

**Chart icon** 

Click this button to generate charts about the used file system storage space for a computer monitoring group or computer reporting group.

**Magnifying glass icon** 

Click this button to view more detailed information about the file systems in a computer monitoring group or computer reporting group.

**Creator**

ID of the user who created a group.

**Group** Name of a computer monitoring group or computer reporting group.

**Consumed File System Space (Used Space)**

Total amount of used file system storage space within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on hypervisors that are assigned to virtual machine servers where the Storage Resource agent or Data agent is installed if both the virtual machine file systems and hypervisor file systems reside within the group.
- File systems on virtual machine servers where the Storage Resource agent or Data agent is installed, but the hypervisor for the virtual machine servers has not been discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems on virtual machine servers that are in a group without the underlying hypervisor file systems, and the total used space exceeds the capacity of the total file system capacity.
- File systems that become missing after a probe.
- File systems that are not mounted.

**Note:** On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

**Percent Used Space**

Percentage of total file system storage space that is used within a group.

**Available File System Space (Free Space)**

Total amount of unique file system storage space that is not used or is available within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are not mounted.

#### **File System Space (Capacity)**

Amount of file system storage space for the computers in a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on virtual machines where the Storage Resource agent or Data agent is installed.
- File systems that become missing after a probe.
- File systems that are not mounted.

**Note:** The value in this column might be greater than **Used Space + Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the `df` command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space + Free Space**, the difference representing the space wasted due to the system overhead.

#### **File Count**

Total number of files within a group. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Note:** For groups that contain Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

#### **Directory Count**

Total number of directories within a group. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Storage Resource agent or Data agent.

#### **Percent Used Inodes**

Percentage of used Inodes within a group.

#### **Used Inodes**

Number of used Inodes within a group.

#### **Free Inodes**

Number of free Inodes within a group.

## By Domain

### Chart icon

Click this button to generate charts about the file system used space for a domain.

### Magnifying glass icon

Click this button to view more detailed information about the file system used space of computers that are in a domain.

### Consumed File System Space (Domain)

Name of a domain.

### Used Space

Total amount of unique file system storage space that is used or consumed by the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**Note:** On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

### Percent Used Space

Percentage of total file system storage space that is used within a domain.

### Available File System Space (Free Space)

Total amount of unique file system storage space that is not used or is available to the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server
- File systems that are not mounted.

### File System Space (Capacity)

Total amount of unique file system storage space that was discovered on the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**Note:** The value in this column might be greater than **Used Space + Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the `df` command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space + Free Space**, the difference representing the space wasted due to the system overhead.

### File Count

Total number of files within a domain. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Note:** On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

### Directory Count

Total number of directories within a domain. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Storage Resource agent or Data agent.

### Percent Used Inodes

Percentage of used Inodes within a domain.

### Used Inodes

Number of used Inodes within a domain.

### Free Inodes

Number of free Inodes within a domain.

### Network-wide

#### Chart icon

Click this button to generate charts about the file system used space in a network.

#### Magnifying glass icon

Click this button to view more detailed information about the file system used space in a network.

### Consumed File System Space (Used Space)

Total amount of unique file system storage space that is used or consumed by the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**Note:** On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

### **Percent Used Space**

Percentage of total file system storage space that is used within a network.

### **Available File System Space (Free Space)**

Total amount of unique file system storage space that is not used or is available to the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server
- File systems that are not mounted.

### **File System Space (Capacity)**

Total amount of unique file system storage space that was discovered on the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**Note:** The value in this column might be greater than **Used Space + Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the `df` command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space + Free Space**, the difference representing the space wasted due to the system overhead.

### **File Count**

Total number of files within a network. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Note:** On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

### **Directory Count**

Total number of directories within a network. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Percent Used Inodes**

Percentage of used Inodes within a network.

**Used Inodes**

Number of used Inodes within a network.

**Free Inodes**

Number of free Inodes within a network.

**TPC-Wide Storage Space - Available File System Space report:**

Use these reports to view detailed information about the available storage space in file systems on systems in your environment. You can view this information in different ways using the following By clauses: by filesystem, by filesystem group, by cluster, by computer, by computer group, by domain, and for the entire network.

**Note:** The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

**By Filesystem****Chart icon** 

Click this button to generate a chart about the used space information for a file system.

**Computer**

Name of the computer on which a file system resides.

**Mount Point**

Name or mount point (UNIX/Linux) of a file system (for example, c:\, d;\, /opt, /export/home, etc.). The file systems displayed in this report include the following:

- File systems on virtual machines that are monitored by IBM Tivoli Storage Productivity Center.
- File systems on hypervisors that are monitored by Tivoli Storage Productivity Center.
- File systems that become missing after a probe.

**Available File System Space (Free Space)**

Amount of unused storage space within a file system.

**Detectable**

Indicates **Yes** when a file system is identified as being available to the server following a successful probe.

**Percent Available File System Space (Percent Free Space)**

Percentage of file system storage space that is unused within a file system.

**Consumed File System Space (Used Space)**

Amount of used storage space within a file system. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

The value that appears in the **TOTAL** row for this column does not include the following file systems:

- File systems on discovered hypervisors that are assigned to virtual machine servers where the Storage Resource agent or Data agent has been installed.
- File systems on discovered virtual machine servers where the Storage Resource agent or Data agent is installed, but the hypervisor for the virtual machine servers has not been discovered and the total used space exceeds the capacity of the total file system capacity.

### **File System Space (Capacity)**

Amount of file system storage space for a file system.

#### **Note:**

- The value that appears in the **TOTAL** row for this column does not include file systems on virtual machines where the Storage Resource agent or Data agent is installed.
- The value in this column might be greater than **Used Space + Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the `df` command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space + Free Space**, the difference representing the space wasted due to the system overhead.

### **File Count**

Total number of files within a file system. This value does not include files on file systems that have not been scanned.

**Note:** On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

### **Directory Count**

Total number of directories within a file system. This value does not include directories on file systems that have not been scanned.

### **Percent Used Inodes**

Percentage of used Inodes within a file system.

### **Used Inodes**

Number of used Inodes within a file system.

### **Free Inodes**

Number of free Inodes within a file system.

### **Volume Group**

Volume group to which a logical volume belongs. If the volume is not part of a volume group this column displays **N/A**.

### **Logical Volume Space (Allocation)**

Logical volume space used by the logical volume containing the file system.

### **Disk Detectable**

Indicates **Yes** when a file system's logical volume is identified as having an underlying host disk.

## By Filesystem Group

### Chart icon

Click this button to generate charts about the used file system storage space in a file system monitoring group or file system reporting group.

### Magnifying glass icon

Click this button to view more detailed information about the used file system storage space in a file system monitoring group or file system reporting group.

### Creator

ID of the user who created a group.

**Group** Name of a file system monitoring group or file system reporting group.

### Available File System Space (Free Space)

Total amount of unique file system storage space that is not used or is available within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are not mounted.

### Percent Available File System Space (Percent Free Space)

Percentage of total file system storage space that is not used or available within a group.

### Consumed File System Space (Used Space)

Total amount of used file system storage space within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on hypervisors that are assigned to virtual machine servers where the Storage Resource agent or Data agent is installed if both the virtual machine file systems and hypervisor file systems reside within the group.
- File systems on virtual machine servers where the Data agents installed, but the hypervisor for the virtual machine servers has not been discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems on virtual machine servers that are in a group without the underlying hypervisor file systems, and the total used space exceeds the capacity of the total file system capacity.
- File systems that become missing after a probe.
- File systems that are not mounted.

**Note:** On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

### File System Space (Capacity)

Amount of file system storage space for the file systems in a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on virtual machines where the Storage Resource agent or Data agent is installed.
- File systems that become missing after a probe.



- File systems that are not mounted.

**Note:** The value in this column might be greater than **Used Space + Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the `df` command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space + Free Space**, the difference representing the space wasted due to the system overhead.

#### File Count

Total number of files within a group. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Note:** On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

#### Directory Count

Total number of directories within a group. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Storage Resource agent or Data agent.

#### Percent Used Inodes

Percentage of used Inodes within a group.

#### Used Inodes

Number of used Inodes within a group.

#### Free Inodes

Number of free Inodes within a group.

#### By Cluster

#### Chart icon

Click this button to generate a chart about the used file system space for a cluster.

#### Magnifying glass icon

Click this button to view more detailed information about the used storage space on file systems that are in a cluster.

#### Cluster

Name of cluster.

### **Available File System Space (Free Space)**

Total amount of unique file system storage space that is not used or is available in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

### **Percent Available File System Space (Percent Free Space)**

Percentage of total file system storage space that is available within a cluster.

### **Consumed File System Space (Used Space)**

Total amount of unique file system storage space that is used or consumed by the file systems in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**Note:** On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

### **File System Space (Capacity)**

Total amount of unique file system storage space in a cluster. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**Note:** The value in this column might be greater than **Used Space + Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the `df` command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space + Free Space**, the difference representing the space wasted due to the system overhead.

### **File Count**

Total number of files within a cluster. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by an agent.

When monitoring computers where the Windows operating system is installed, the number of files indicated in this column might be lower than the number of files reported by the operating system.

**Directory Count**

Total number of directories within a cluster. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by an agent.

When monitoring computers where the Windows operating system is installed, the number of folders indicated in this column might be higher than the number of folders that is reported by the operating system.

**Percent Used Inodes**

Percentage of used Inodes within a cluster.

**Used Inodes**


Number of used Inodes within a cluster.

**Free Inodes**

Number of free Inodes within a cluster.

**By Computer****Chart icon** 

Click this button to generate a chart about the used file system storage space of the computers that are displayed in the report.

**Magnifying glass icon** 

Click this button to view a more detailed report about the used file system storage space on a computer.

**Computer**

Name of the computer on which a file system resides.

**Available File System Space (Free Space)**

Total amount of unique file system storage space that is not used or is available to a computer. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**Percent Available File System Space (Percent Free Space)**

Percentage of total file system storage space that is available within a computer.

**Consumed File System Space (Used Space)**

Total amount of unique file system storage space that is used or consumed by the computer where the Storage Resource agent or Data agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**Note:** On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

### File System Space (Capacity)

Total amount of unique file system storage space that was discovered on the computer where the Storage Resource agent or Data agent is installed. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

#### **Note:**

- The value in this column might be greater than **Used Space + Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the `df` command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space + Free Space**, the difference representing the space wasted due to the system overhead.
- The value that appears in the **TOTAL** row for this column does not include file systems on virtual machines where the Storage Resource agent or Data agent is installed.
- The drill-down reports for each computer will generate a list of detectable file systems only. Therefore, the total capacity of file systems for each computer drill-down report will equal this value.

### File Count

Total number of files within a computer. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Note:** On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

### Directory Count

Total number of directories within a computer. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Percent Used Inodes**

Percentage of used Inodes within a computer.

**Used Inodes**

Number of used Inodes within a computer.

**Free Inodes**

Number of free Inodes within a computer.

**By Computer Group****Chart icon** 

Click this button to generate charts about the used file system storage space for a computer monitoring group or computer reporting group.

**Magnifying glass icon** 

Click this button to view more detailed information about the file systems in a computer monitoring group or computer reporting group.

**Creator**

ID of the user who created a group.

**Group** Name of a computer monitoring group or computer reporting group.

**Available File System Space (Free Space)**

Total amount of unique file system storage space that is not used or is available within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are not mounted.

**Percent Available File System Space (Percent Free Space)**

Percentage of total file system storage space that is available within a group.

**Consumed File System Space (Used Space)**

Total amount of used file system storage space within a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on hypervisors that are assigned to virtual machine servers where the Storage Resource agent or Data agent is installed if both the virtual machine file systems and hypervisor file systems reside within the group.
- File systems on virtual machine servers where the Storage Resource agent or Data agent is installed, but the hypervisor for the virtual machine servers has not been discovered and the total used space exceeds the capacity of the total file system capacity.
- File systems on virtual machine servers that are in a group without the underlying hypervisor file systems, and the total used space exceeds the capacity of the total file system capacity.
- File systems that become missing after a probe.
- File systems that are not mounted.

**Note:** On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

### File System Space (Capacity)

Amount of file system storage space for the computers in a group. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems on virtual machines where the Storage Resource agent or Data agent is installed.
- File systems that become missing after a probe.
- File systems that are not mounted.

**Note:** The value in this column might be greater than **Used Space + Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the `df` command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space + Free Space**, the difference representing the space wasted due to the system overhead.

### File Count

Total number of files within a group. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Note:** For groups that contain Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

### Directory Count

Total number of directories within a group. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Storage Resource agent or Data agent.

### Percent Used Inodes

Percentage of used Inodes within a group.

### Used Inodes

Number of used Inodes within a group.

### Free Inodes

Number of free Inodes within a group.

### By Domain

### Chart icon

Click this button to generate charts about the file system used space for a domain.

### Magnifying glass icon

Click this button to view more detailed information about the file system used space of computers that are in a domain.

### Domain

Name of a domain.

### Available File System Space (Free Space)

Total amount of unique file system storage space that is not used or is available to the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server
- File systems that are not mounted.

### Percent Available File System Space (Percent Free Space)

Percentage of total file system storage space that is available within a domain.

### Consumed File System Space (Used Space)

Total amount of unique file system storage space that is used or consumed by the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**Note:** On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

### File System Space (Capacity)

Total amount of unique file system storage space that was discovered on the computers within a domain. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**Note:** The value in this column might be greater than **Used Space + Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the `df` command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space + Free Space**, the difference representing the space wasted due to the system overhead.

### File Count

Total number of files within a domain. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Note:** On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

#### **Directory Count**

Total number of directories within a domain. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Storage Resource agent or Data agent.

#### **Percent Used Inodes**

Percentage of used Inodes within a domain.

#### **Used Inodes**

Number of used Inodes within a domain.

#### **Free Inodes**

Number of free Inodes within a domain.

#### **Network-wide**

##### **Chart icon**

Click this button to generate a chart about the file system used space in a network.

##### **Magnifying glass icon**

Click this button to view more detailed information about the file system used space in a network.

#### **Available File System Space (Free Space)**

Total amount of unique file system storage space that is not used or is available to the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server
- File systems that are not mounted.

#### **Percent Available File System Space (Percent Free Space)**

Percentage of total file system storage space that is available within a network.

#### **Consumed File System Space (Used Space)**

Total amount of unique file system storage space that is used or consumed



by the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**Note:** On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

#### **File System Space (Capacity)**

Total amount of unique file system storage space that was discovered on the computers within a network. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**Note:** The value in this column might be greater than **Used Space + Free Space** on some UNIX systems. This occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the `df` command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **Capacity** field is greater than or equal to the value of **Used Space + Free Space**, the difference representing the space wasted due to the system overhead.

#### **File Count**

Total number of files within a network. Tivoli Storage Productivity Center does not include the following files in its calculation of the value for this column:

- Files on file systems that become missing after a probe.
- Files on file systems that are not mounted.
- Files on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Note:** On Windows computers, AIX systems, and systems using IBM General Parallel File System (GPFS) for file management, IBM Tivoli Storage Productivity Center calculates both the number of files and the number of directories for the value in this column. For example, if a Windows computer contains 150 files and 50 directories, the number in the File Count column is 200.

#### **Directory Count**

Total number of directories within a network. Tivoli Storage Productivity Center does not include the following directories in its calculation of the value for this column:

- Directories on file systems that become missing after a probe.
- Directories on file systems that are not mounted.
- Directories on file systems that have not been scanned by a Storage Resource agent or Data agent.

**Percent Used Inodes**

Percentage of used Inodes within a network.

**Used Inodes**

Number of used Inodes within a network.

**Free Inodes**

Number of free Inodes within a network.

**Usage Reports**

Use Usage reports to view the usage and growth of the consumption of your storage. You can view these reports based on specific filesystems and computers, groups of filesystems and computers, or throughout your entire enterprise. These reports help you project storage consumption for the future and maximize the current storage assets that you have in place by eliminating wasted space and making the most of the space you have.

Use these reports to do the following:

- View which servers and filesystems are experiencing the heaviest (or lightest) load of data and storage access.
- Identify wasted space by pinpointing files that are no longer needed or have not been accessed for the longest time.
- Associate users with the amount of storage they are consuming on specific filesystems, groups of filesystems, specific computers, groups of computers, and throughout the entire network.
- Select the files against which to run an IBM Tivoli Storage Manager archive/backup job. Tivoli Storage Manager protects your organization's data from hardware failures and other errors by storing backup and archive copies of data on offline storage.

**Access Load reports:** Use Access Load reports to view which machines and file systems within your environment are experiencing the heaviest (or lightest) load of data and storage access.

**Access Time reports**

Use Access Time reports to view and chart the amount of data and the number of files that have been accessed during the last day, the last week, the last month, the last year, and over. You can view this information at the directory level, the file system level, the computer level, at the domain level, and for the entire network.

The following report shows a single pie chart that depicts all of the data across the entire enterprise, giving you an easy way to see how much of your corporate data is accessed every day, every week, and even data that has not been accessed in over a year.

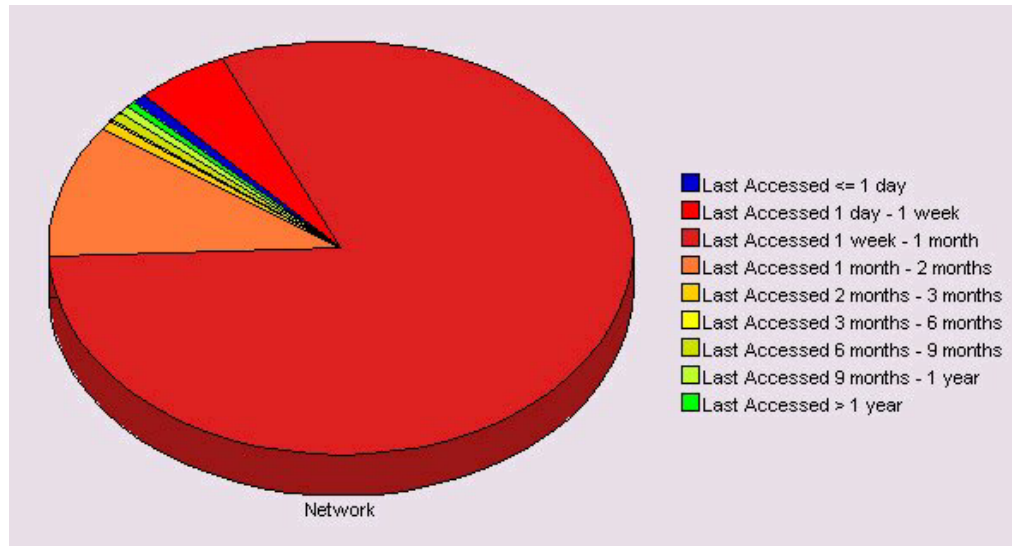


Figure 58. Usage Reporting—Access Load—Access Time—Chart Network Wide

### Modification Time reports

Use Modification Time reports to view and chart the amount of data and the number of files that have been modified during the last day, the last week, the last month, the last year, and over. You can view this information at the directory level, the file system level, the computer level, at the domain level, and for the entire network.

The following report shows a single pie chart that depicts all of the data across the entire enterprise, allowing you to see how much of your corporate data is modified every day, every week, and even data that has not been modified in over a year.




Figure 59. Usage Reporting—Access Load—Modification Time—Chart Network Wide

### Creation Time reports

Use Creation Time reports to view and chart the amount of data and the number of files that have been created during the last day, the last week,

the last month, the last year, and over. You can view this information at the directory level, the file system level, the computer level, at the domain level, and for the entire network.

**Note:** Because scans do not collect creation time information for the files stored on UNIX machines, creation time reports do not contain any data for the non-Windows machines in your environment.

**Files reports:** Use Files Reports to report on the files found on the storage in your environment during the Data Manager scan process. During the scan process, Data Manager gathers statistics on your files and attributes about the files. Select the files against which to run an IBM Tivoli Storage Manager archive or backup job. See “Archive and backup functions” on page 498 for information on how to create archive and backup jobs from the files shown in a Files report. The information gathered can be reported on in a variety of ways using this facility:

#### **Largest Files reports**

Use these reports to view information about the largest files in your environment. Largest Files reports enable you to do the following:

- Generate reports that contain detailed information about the largest files found in your environment. You can view reports on the largest files in the following ways: by directory, by directory group, by file system, by file system group, by cluster, by computer, by computer group, by domain, and for the entire network.

**Note:** The number of files that are gathered and stored in the repository is dependent upon the number that you specify in the profile that you use during scanning. The default profile gathers information on the 20 largest files. You can select a profile for the report by clicking the **Profile** list box.

- Select files against which to run an Tivoli Storage Manager archive or backup job. See “Archive and backup functions” on page 498 for information on how to create archive/backup jobs from the files shown in a Files report.
- Add files to an existing Tivoli Storage Manager archive/backup job. See “Archive and backup functions” on page 498 for information on how to add files from reports to an archive/backup job.

#### **Duplicate Files reports**

Use these reports to view information about the files found in your environment during a scan that contain duplicate file names. Duplicate Files reports enable you to do the following:

- Generate reports on the files found in your environment during a scan that contain duplicate file names. Use this report to view a list of files that might no longer be needed and could be wasting storage space. You can use the data generated in a duplicate files report to view a list of files that might no longer be needed and could be wasting storage space. You can view reports on duplicate files in the following ways: by directory, by directory group, by file system, by file system group, by cluster, by computer, by computer group, by domain, and for the entire network.

**Note:** The number of files that are gathered and stored in the repository is dependent upon the number that you specify in the profile that you

use during scanning. The default profile gathers information on the 20 largest files. You can select a profile for the report by clicking the **Profile** list box.

- Select files against which to run an Tivoli Storage Manager archive or backup job. See “Archive and backup functions” on page 498 for information on how to create archive and backup jobs from the files shown in a Files report.
- Add files to an existing Tivoli Storage Manager archive or backup job. See “Archive and backup functions” on page 498 for information on how to add files from reports to an archive or backup job.

Before generating a Duplicate Files report for environments that contain a large amount of files (i.e., environments that might contain over 100,000 duplicate files), we recommend you tune the Data Manager repository. To do this, run the appropriate RDBMS utility to collect statistics on the repository tables. Run the following utilities depending on the RDBMS where the repository is stored:

- UDB: runstats
- Oracle: analyze
- Microsoft SQL Server: update stats
- Sybase: update stats

Run the appropriate utility against all the tables in the repository schema you specified during the installation process.

To generate the data for a duplicate files report, you can either create a constraint that will save filenames for *n* number of violation files, or run a scan that includes one or all of the following existing profiles:

- TPCUser.Largest Files
- TPCUser.Largest Orphans
- TPCUser.Most at Risk
- TPCUser.Most Obsolete

There is not a specific Duplicate Files profile that collects information on all duplicate files in a file system. The filenames that are triggered by the previously listed profiles or constraints will be saved to the repository and the duplicate files report is generated from the data gathered from the scan or constraint.

For example, create a scan using the Largest Files profile to collect information on the 100 largest files. The 100 largest files are saved in the `t_stat_file` table in the repository. If five duplicate files were saved as a result of being flagged by the Largest File Profile, then those files will be reported as duplicates in the Duplicate file Report.

### **Most Obsolete Files reports**

Use these reports to view information about the files found in your environment that have not been accessed in the longest period of time. Most Obsolete Files reports enable you to do the following:

- Generate reports on the files found in your environment that have not been accessed in the longest period of time. It is very likely that many of these files are no longer needed, at least online, and could be wasting space. You can view reports on the largest files in the following ways: by directory, by directory group, by file system, by file system group, by cluster, by computer, by computer group, by domain, and for the entire network.

**Note:** The number of files that are gathered and stored in the repository is dependent upon the number that you specify in the profile that you use during scanning. The default profile gathers information on the 20 largest files. You can select a profile for the report by clicking the **Profile** list box.

- Select files against which to run a Tivoli Storage Manager archive or backup job. See “Archive and backup functions” on page 498 for information on how to create archive/backup jobs from the files shown in a Files report.
- Add files to an existing Tivoli Storage Manager archive or backup job. See “Archive and backup functions” on page 498 for information on how to add files from reports to an archive or backup job.

### Orphan Files reports

Use these reports to view information about files that are owned by users that are no longer active in the environment. Orphan files reporting enables you to do the following:

- Generate reports on files that are owned by users that are no longer active in the environment. It is very likely that many of these files are no longer needed and could be wasting space. You can view reports on orphan files in the following ways: by directory, by directory group, by file system, by file system group, by cluster, by computer, by computer group, by domain, and for the entire network.

You can also generate reports on the oldest orphan files that have not been accessed for the longest periods of time, as well as largest orphan files.

**Note:** The number of files that are gathered and stored in the repository is dependent upon the number that you specify in the profile that you use during scanning. The default profile gathers information on the 20 largest files. You can select a profile for the report by clicking the **Profile** list box.

- Select files against which to run a Tivoli Storage Manager archive or backup job. See “Archive and backup functions” on page 498 for information on how to create archive and backup jobs from the files shown in a Files report.
- Add files to an existing Tivoli Storage Manager archive or backup job. See “Archive and backup functions” on page 498 for information on how to add files from reports to an archive or backup job.

### File Size Distribution reports

Use these reports to view information about the distribution of file sizes across storage resources. Use File Size Distribution reports to generate reports that show the distribution of file sizes on a specific directory, a file system, a computer, the domain, or across your entire network.

The following report shows a single pie chart for each computer in the report.

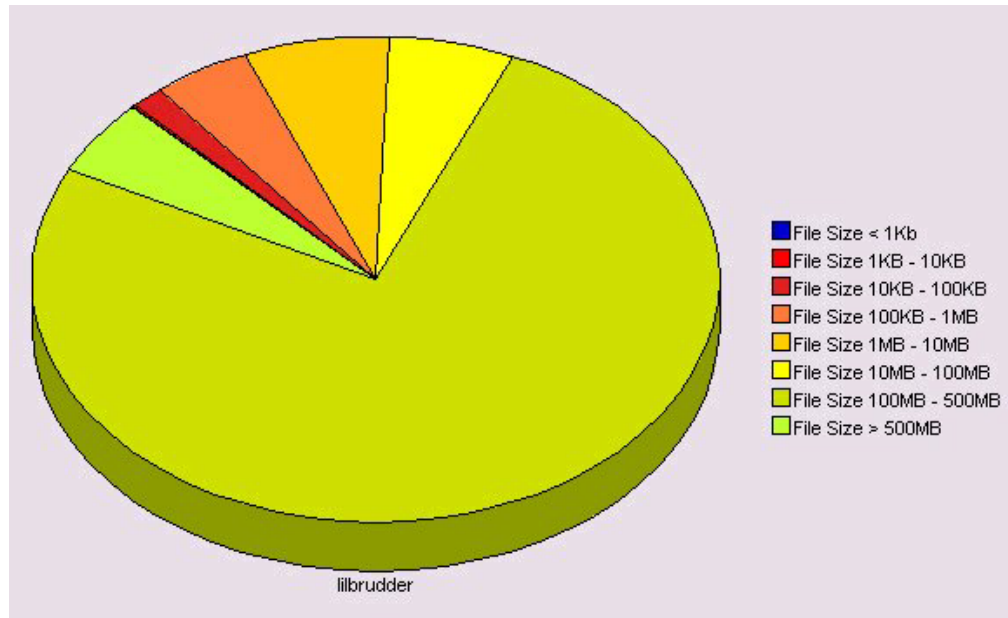


Figure 60. Usage Reporting—Files—File Size Distribution—File Count Chart Tab

Each pie chart depicts the number of files that fall into each file size bucket that are described in the legend. For example, the blue portion of the pie represents the number of files on each computer that are smaller than 1024 bytes long. This gives you an easy way to get a quick grasp of what are the size distribution of your data files.

### File Summary reports

Use these reports to view general information about the files in your environment. Use File Summary reports generate overview information for files, also at the directory level, the file system level, the computer level, the domain level or for the whole network.

This report shows you the following (for each entity at the level you are reporting by):

- The total size of the files
- The number of files
- The number of directories
- The average file size

### File Types reports

Use these reports to view information about the storage usage of different file types in your environment. File Types reports enable you to generate reports on the file types found in your environment during a scan. The scan that collects data for these reports must use the TPCUser.Summary By File Type profile. For example, when you run a scan using the TPCUser.Summary By File Type profile, Data Manager will collect storage usage information about such file types as .exe, .zip, .sys, .pdf, .doc, .dll, .wav, .mp3, .avi, etc. Each of these file types is represented by its own row in the generated reports.

Use these reports to do the following:

- Relate the space used by applications to the total capacity and used space. For example, you can use these reports to view the total amount of storage consumed by Adobe Acrobat files and Lotus Notes mail databases.
- View which applications are consuming the most space across a given set of storage resources

You can view these reports in the following ways:

- **Group by Resource.** These reports organize data by storage resource (for example, file system, computer, cluster, etc.) and space used.
- **Group by File Type.** These reports organizes data by file type and space used. They show where the most spaced used by each type of file resides.

### Archive/Backup jobs from reports:

Use the Archive/Backup function from Reporting to define an archive or backup job to run against the files that you select from specified reports.

Tivoli Storage Manager protects your organization's data from hardware failures and other let you back up specified files. These files are not excluded in the include-exclude list and meet the requirement for serialization in the backup copy group of the management class assigned to each file. An incremental backup enables you to back up new or changed files or directories from a client domain or from specified directories or files. These directories or files are not excluded in the include-exclude list and meet the requirements for frequency, mode, and serialization as defined by a backup copy group of the management class assigned to each file. See *IBM Tivoli Storage Manager Client User's Guide* for more information about Tivoli Storage Manager archives and backups.

This function enables you to select a specific file or group of files from Data Manager reports that you want to archive or back up using the Tivoli Storage Manager tool.

Archive/backup is run from **Data Manager -> Reporting**. You cannot create archive and backup jobs from the Archive/Backup node - rather you create these jobs by selecting files from the previously listed reports or the reports that you create. Each archive and backup that you create from reports appears as a node under the Archive/Backup node. The reports that contain files against which you can run an archive or backup job include:

- **Data Manager -> Reporting -> Usage -> Files -> Largest Files**
- **Data Manager -> Reporting -> Usage -> Files -> Most Obsolete Files**
- **Data Manager -> Reporting -> Usage -> Files -> Orphan Files**
- **Data Manager -> Reporting -> Usage -> Files -> Duplicate Files**
- **Data Manager -> Reporting -> Backup -> Most at Risk Files ->**
- **Data Manager -> Reporting -> Usage Violations -> Constraint Violations**

You can also run an archive or backup job from a report you create.

### Note:

See *Creating Archive/Backup jobs in Reporting* for information about how to create an archive or backup job from the files shown in reports.



### **Directories reports:**

Use Directories reports to report upon the directories found on the storage in your environment during the scan process. During the scan process, Data Manager gathers a vast number of statistics on your directories and a number of attributes about the directories. Before you can view information about directories in these reports, you must do the following:

1. Use the **Monitoring → Groups → Directory** function to create directory groups that contain the directories upon which you want to report.
2. Include the directory groups you created in scans.
3. Schedule and run the scans that contain the directory groups (and directories) upon which you want to report.

### **Largest Directories reports**

Use Largest Directories reports to view reports on the  $n$  largest directories found in your environment:

- for specific directories
- on specific computers or groups of computers
- on specific file systems or groups of file systems
- in specific domains
- on the entire network

The number of largest (or the value of  $n$  in finding the  $n$  largest directories), is determined by the profile that you use for doing the reporting. When you define a profile in the **Monitoring → Profiles** node, you specify the number of largest directories upon which to collect information. When you generate a Largest Directories report, you can select what profile to use.

### **Monitored Directory Space reports**

Use Monitored Directory reports to report on the monitored directories that you have included in a group via the **Monitoring → Groups → Directory** facility and included in a scan. During the scan process, Data Manager gathers a number of statistics on your directories and a number of attributes about the directories.

**Users reports:** Use Users Reports to view information about the users that are found as file owners in your environment. These reports provide an easy and convenient way to view and aggregate storage resource consumption by owner across directories, file systems, computers, domains, and across your entire network.

**OS User Groups reports:** OS User Groups reports give you the exact same capabilities as the reports under Users reports. The only difference is that you are doing your analysis of OS User Groups as opposed to actual single users.

### **Usage Violation reports**

Use Usage Violation reports to view information about quota and constraint violations.

Usage Violation reports enable you to:

- Create and enforce corporate storage usage policies based on reports that show quota and constraint violations.

Use quotas to control how much storage a user, or a group of users, can consume on a filesystem or group of filesystems, a computer or group of computers, and throughout the entire network. Use Constraints to set policies

about which types of files, owners, and file sizes are allowed on a filesystem or group of file systems, a computer or group of computer, and throughout the entire network.

- Select files in Constraint Violation reports against which to run an IBM Tivoli Storage Manager archive/backup job. See “Archive and backup functions” on page 498 for information on how to create archive/backup jobs from the files shown in a Files report.
- Add files from Constraint Violation reports to an existing Tivoli Storage Manager archive or backup job. See “Create an archive/backup job” on page 499 for information on how to add files from reports to an archive/backup job.

**Constraint Violations reports:** Constraint Violations reporting enables you to do the following:

- View the files that have violated a constraint. The report shows one row for each Constraint. The row shows the detail of the constraint and the violating space totals for the files that violate the constraint. You can view reports on the files that violate constraints in the following ways: by file system, by file system group, by computer, by computer group, by domain, and by constraint.
- Select files against which to run an Tivoli Storage Manager archive or backup job. See the “Archive and backup functions” on page 498 section for information on how to create archive and backup jobs from the files shown in a Files report.
- Add files to an existing Tivoli Storage Manager archive or backup job. See the “Archive and backup functions” on page 498 section for information on how to add files from reports to an archive or backup job.

**Quota Violations reports:** Use Quota Violations reports to report on the quotas that you have defined and see how much space is being used by the various entities that have quotas defined upon them.

Quotas can be defined on either of two categories of users: on a specific user or group of users that you have defined within either Users or on OS User Groups.

| Quota Category | Description                                                                                                                                                                           |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| User           | This allows you to define a quota on a user or group of users. The users are the owners of the files that Data Manager finds as it scans your systems.                                |
| OS User Group  | This allows you to define a quota on an OS user group or group of OS user groups. OS user groups are defined outside the tool and are found by Data Manager as it scans your systems. |

Within each of those categories, you can define quotas at any one of following three levels:

| Quota Level | Description                                                                                                                                                                                                                            |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Network     | This allows you to define quotas on users or groups of users at the entire network level. This enables you to control how much any single user or group of users is using on all of the storage throughout your enterprise.            |
| Computer    | This allows you to define quotas on users or groups of users at the computer level. This enables you to control how much any single user or group of users is using on a single computer, or on any given group of computers.          |
| Filesystem  | This allows you to define quotas on users or groups of users at the file system level. This enables you to control how much any single user or group of users is using on a single file system, or on any given group of file systems. |

## All Quotas reports

Use All Quotas reports to report on all quotas that have been violated, including those at the network level, computer level, and file system level. The **Selection** tab allows you to select upon what users you wish to do quota reporting.

The **Objects** tab shows you for each user and computer and file system:

- how much space is being used
- what the Quota threshold is
- what the percentage of the Quota that is being consumed

This gives you a simple way to analyze how much space is being used relative to the Quotas that you have defined on your systems and users.

## Network-wide Quotas reports

Use Network-wide Quota reports to view all the network-wide quotas that have been violated by users and OS user groups in your environment. Information about the quotas that appear on these reports includes the amount of space being used, the threshold value, and the percentage above the threshold value. You can view these reports sorted by user, by OS user group, and by the creator of the quota.

## Computer Quotas reports

Use Computer Quota reports to view all the computer quotas that have been violated by users and OS user groups in your environment. Information about the quotas that appear on these reports includes the amount of space being used, the threshold value, and the percentage above the threshold value. You can view these reports sorted by user, by OS user group, and by the creator of the quota.

## Filesystem Quotas reports

Use Filesystem Quota reports to view all the computer quotas that have been violated by users and OS user groups in your environment. Information about the quotas that appear on these reports includes the amount of space being used, the threshold value, and the percentage above the threshold value. You can view these reports sorted by user, by OS user group, and by the creator of the quota.

## Backup reports

Use Backup reports to identify the most at risk files in your enterprise that are not backed up properly. Critical data that is not backed up might expose you to costly losses of information.

Backup reports also enable you to:

- Investigate the sizes of data that would be required to backup to achieve full backups of any subset of your enterprise (including the whole enterprise).
- Determine the amount of space required to perform an incremental backup of any or all of your data at any given point in time.
- Select files in Most at Risk Files reports against which to run an IBM Tivoli Storage Manager archive/backup job.
- Add files from Most at Risk Files reports to an existing Tivoli Storage Manager archive/backup job.

**Most At Risk Files reports:** Most At Risk Files reporting enables you to do the following:

- View the files that have the earliest modification time, but have not been backed up since they were modified. This can alert you to a dangerous situation where you may have vulnerable data that could be lost and not recovered due to lack of backups.
- Select files against which to run an IBM Tivoli Storage Manager archive or backup job. See the “Archive and backup functions” on page 498 section for information on how to create archive or backup jobs from the files shown in a Files report.
- Add files to an existing Tivoli Storage Manager archive or backup job. See the “Archive and backup functions” on page 498 section for information on how to add files from reports to an archive or backup job.

**Modified Files not Backed Up reports:** Use Modified Files not Backed Up reports to view the numbers of files that have the earliest modification time, but have not been backed up since they were modified. This can alert you to a dangerous situation where you may have vulnerable data that could be lost and not recovered due to lack of backups.

**Backup Storage Requirements reports:** Use Backup Storage Requirements reports to determine, at a network, domain computer, file system, or directory level, how much storage will be required to do backups. You can determine how much storage will be required to do both full backups and to do incremental backups.

| Reporting Type                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Full Backup Size</b>        | Use Full Backup Size reports to view how much space it will take to do full backups of specific computers, file systems, directories, groups of those entities, domains, and the entire network.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Incremental Backup Size</b> | (Windows only) Use this report to view the total size of all the files that have not been backed up since the last scan. You can generate this report by Directory, by Directory Group, by Filesystem, by Filesystem Group, by Computer, by Computer Group, by Domain and Network Wide. This allows you to see for any of these levels, how much data would need to be backed up for any of these desired time intervals. This is useful for backup sizing and can be used to determine appropriate backup groups and sets for timing your backups.                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Incremental Range Sizes</b> | Use this report to determine how much storage space you would need to back up data that has been modified in the last day, week, month, two months, three months, six months, nine months, or year. For example, if you want to back up the files that have been modified in the last 3 months only, generate this report and view the values in the "Last Modified <= 3 months" column to determine 1) how many files changed and 2) how much storage those changed files are consuming.<br><br>You can generate this report by Directory, by Directory Group, by Filesystem, by Filesystem Group, by Cluster, by Computer, by Computer Group, by Domain and Network Wide. This allows you to see for any of these levels, how much data would need to be backed up for any of these desired time intervals. This is useful for backup sizing and can be used to determine appropriate backup groups and sets for timing your backups. |

## Monitored Computer Storage Space reports

Use these host-based reports to view the data that is collected by probes for computers, network attached storage, and hypervisors. These reports contain detailed information about the consumption of file system space and enable you to relate that space back to external or internal storage without counting shared space multiple times.

Monitored computer storage space reports are organized into the following categories:

### Disk Storage

Use these reports to view information about disks that reside on computers monitored by IBM Tivoli Storage Productivity Center. This includes information about storage subsystem disks, known disks with serial numbers, or unknown disks without serial numbers.

### Non-Disk Storage

Use these reports to view information about storage space that does not reside on computer disks that are monitored by Tivoli Storage Productivity Center. This type of storage includes the following:

- Storage that resides on a remote mount or a network attached server.
- Storage that is temporarily allocated to processor memory.

You must perform the following actions before viewing these reports:

- Run a probe against the storage entities upon which you want to report.
- If you want to generate By Computer Group, By Filesystem Group, and By Storage Subsystem Group reports, perform the following tasks:
  - Create computer monitoring group and file system monitoring groups using the **Data Manager > Monitoring > Groups** node and then probe them.
  - Create storage subsystem monitoring group using the **Disk Manager > Monitoring > Groups** node and probe them.
  - Create a computer reporting group or file system reporting group using the **Data Manager > Reporting > Groups** and **Disk Manager > Reporting > Groups** node. Probe the entities in the group.
  - Create a subsystem reporting group using the **Disk Manager > Reporting > Groups** node. Probe the storage subsystems in the group.

### Related Topics

- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

### Monitored Computer Storage Space reports - Storage entities:

Monitored Computer Storage Space reports display information about storage entities that you probe. Use this topic to view descriptions of those entities.

You can use Monitored Computer Storage Space reports to view information about the following storage entities:

#### Computer

Computer that is monitored by a Storage Resource agent or Data agent. This computer can be a single host, a virtual machine, or a hypervisor. A hypervisor is a machine that hosts one or more virtual machines.

**Disk on a storage subsystem**

Computer disk that originates from a storage subsystem. If the disk's storage subsystem is not monitored by IBM Tivoli Storage Productivity Center, that storage subsystem is represented by a portion of the serial number of the disk. If the storage subsystem is monitored by Tivoli Storage Productivity Center, that storage subsystem is represented by the name collected by a probe.

**Disk not on a storage subsystem with a serial number**

Computer disk that Tivoli Storage Productivity Center is unable to determine if it originates from a storage subsystem. Tivoli Storage Productivity Center is able to extract the serial number of that disk.

**Disk not on a storage subsystem without a serial number**

A valid computer disk for which Tivoli Storage Productivity Center is unable to do the following:

- Determine if the disk originates from a storage subsystem.
- Extract the serial number of the disk.

**Disk Group**

A collection of disk drives on a computer.

**File System**

Storage entity that creates, removes, and modifies files on a computer. File systems can be shared among monitored computers and span portions of different disks. If a file system spans disk entities, then there will be a single logical volume beneath it that describes how the file system is spanned.

**(Parent) Logical Volume**

Logical representation of a volume that spans many logical volumes. If there is no file system created on this logical volume, information about the parent logical volume is displayed in the reports.

**(Child) Logical Volume**

Logical representation of a volume that spans one or more portions of a single or multiple disks. This entity helps determine how the parent logical volumes and file systems are allocated on the disks. If there is no file system created on this logical volume, or there is no parent logical volume for this logical volume, the child logical volume is displayed in the reports.

**Non-Disk Storage**

Storage space that does not reside on a computer disk. This includes storage residing on a remote mount or a network attached server and storage that is temporarily allocated to processor memory.

**Volume Group**

A collection of logical volumes on a computer.

**Virtual Machine (VM) Disk File**

Storage entity that represents a computer disk on a computer that is a virtual machine.

**Virtual Machine File System (VMFS)**

Storage entity that creates, removes, and modifies VM Disk files on a computer that is a hypervisor hosting one or more virtual machines.

**Monitored Computer Storage Space reports - Common storage space values:**

Monitored Computer Storage Space reports provide detailed storage space information about the computers probed by IBM Tivoli Storage Productivity Center. Use this topic to view descriptions of the common storage values that appear in these reports.

You can view the following storage space information in Monitored Computer Storage Space reports:

**Space** Total amount of storage space for the entity that a row represents:

- For file system rows, this is the total logical volume space where a file system resides.
- For logical volume rows, this is the total logical volume space.
- For computer disk rows, this is the total disk space of a computer disk.
- For computer space rows, this is the total storage subsystem volume space on a computer.
- For storage subsystem space rows, this is the total storage subsystem volume space on all monitored computers.

### **Consumed Space**

Amount of consumed storage space for the entity that a row represents:

- For file system rows, this is the space that is consumed by data that can be accessed by the end user of a computer.
- For logical volume rows, this is the space that can be consumed by a file system, but is not yet consumed by a file system.
- For computer disk rows, this is the space that is consumed by the file systems and logical volumes on a disk.
- For computer space rows, this is the storage subsystem volume space that is consumed on a monitored computer.
- For storage subsystem space rows, this is the storage subsystem volume space that is consumed on all monitored computers.

### **Available Space**

Amount of available storage space for the entity that a row represents:

- For file system rows, this is the space that is available for consumption.
- For logical volume rows, this value is zero (0) because no file system has been created on the logical volume that would help determine the amount of available space.
- For computer disk rows, this is the space that is available on the file systems on a computer disk and the space on a computer disk not yet allocated to a logical volume.
- For computer space rows, this is the storage subsystem volume space that is available for consumption on a computer.
- For storage subsystem space rows, this is the storage subsystem volume space that is available for consumption on all monitored computers.

### **Unavailable Space**

Amount of unavailable storage space for the entity that a row represents:

- For file system rows, this space is unavailable for consumption because it is reserved by the software RAID level on the underlying logical volume and is required for data that is used by the operating system to maintain a file system.
- For logical volume rows, this space is unavailable because it is reserved for the software RAID level of a logical volume.

- For computer disk rows, this is the space that is unavailable on the file systems and the logical volumes on a computer disk.
- For computer space rows, this is the storage subsystem volume space that is unavailable on a computer.
- For storage subsystem space rows, this is the storage subsystem volume space that is unavailable on all monitored computers.

### **Monitored Computer Storage Space - Disk Storage reports:**

Use these reports to view information about disks residing within computers that are monitored by IBM Tivoli Storage Productivity Center. This includes information about storage subsystem disks, known disks with serial numbers, and unknown disks without serial numbers.

Disk Storage reports are organized into the following categories:

#### **All Disk Storage**

Use these reports to view information about all the disks residing on computers that are monitored by Tivoli Storage Productivity Center.

#### **Disk Storage On Storage Subsystems**

Use these reports to view information about disks that Tivoli Storage Productivity Center has identified as storage subsystem volumes.

#### **Disk Storage Not On Storage Subsystems**

Use these reports to view information about disks that Tivoli Storage Productivity Center has not identified as storage subsystem volumes.

#### **Related Topics**

- Monitored Computer Storage Space reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

#### *Disk Storage - All Disk Storage reports:*

Use these reports to view information about all the disks residing on computers that are monitored by IBM Tivoli Storage Productivity Center.

See All Disk Storage reports - Column Descriptions for descriptions of all the columns that appear in All Disk Storage reports.

Use By clauses to filter the entities that appear in a report. For example, if you want to view an All Disk Storage report that displays specific computers, select **By Computer** and use the **Selection** and **Filter** options to determine which computers appear in that report.

You can customize the information in each report by relating the monitored entities to other storage. The following options for relating entities to storage in these reports includes the following:

- Use the **Relate To:** options in the By Computer, By Computer Group, and By Disk/Volume Group reports to relate storage entities to the following storage:
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space



- Computer Disk Space
- Distributed Computer Disk Space
- Use the **Relate To:** options in the By File System/Logical Volume and By File System Groups reports to relate storage entities to the following storage:
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space

### Related Topics

- All Disk Storage reports - Column Descriptions
- Monitored Computer Storage Space reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

#### *All Disk Storage reports - Column Descriptions:*

Use these reports to view information about all the disks residing on computers that are monitored by IBM Tivoli Storage Productivity Center. Use this topic to view descriptions of the columns that appear in these reports.

The following buttons and columns appear in the By clauses of All Disk Storage reports:

**Note:** Not all of the following columns appear in each By clause.

#### **Selection...**

Click this button to select the specific entities that you want to include in the report. For example, in the **By Filesystem/Logical Volume** report, click this button to select the specific file systems or logical volumes upon which you want to report.

**Filter** Click this button to apply conditions to the columns within a report to further select or restrict the entities that you want to include in the report.

#### **Relate <entity> To:**

Select how you want to relate monitored storage entities in the report to other storage.

#### **Magnifying glass icon**

This button appears on all **By <entity> Group** reports. Click this button to view a more detailed report about the entities in a group.

#### **Available Disk Space**

Amount of space on a disk that is available for use. For computer groups, this is the amount of space on disks that is available for use by the computers in a group. For disk and volume groups, this is the external disk space available for use within those groups.

This space includes the file system space that is available on the disk and the space that has not been allocated to a logical volume.

**Note:** The XIV system GUI and the XIV system command-line interface (CLI) report sizes in GB of disks, volumes, pools, and so on, based on  $10^9$ . Tivoli Storage Productivity Center reports sizes in GB based on  $2^{30}$ . For example, a volume showing a size of 17 GB in the XIV system GUI is

reported as 16 GB (XIV GB size \* 1000 \* 1000 \* 1000 / 1024 / 1024 / 1024 rounded to the nearest GB) in Tivoli Storage Productivity Center.

**Available Disk/Volume Group Space**

Amount of available space in a disk group or volume group.

**Available File System/Logical Volume Space**

Amount of space on the file system or logical volume that is available for use. For computer groups, this is the amount of space on file systems or logical volumes that is available for use by the computers within those groups. For file system groups, this is the amount of space on file systems or logical volumes within those groups that can be consumed. For disk groups and volume groups, this is the external logical volume or file system space available for use within those groups.

**Computer**

Name of a computer where a disk, file system, logical volume, or disk group is located.

**Computer OS Type**

The operating system running on a computer.

**Consumed Disk Space**

Amount of space on a disk that is consumed. For computer groups, this is the amount of disk consumed by the computers within those groups. For disk and volume groups, this is the external disk space consumed within those groups.

This space includes the file system space that is consumed on a disk. If a logical volume on this disk space has no file system, then this space includes the consumed portion of the logical volume that has been allocated to the disk.

**Consumed Disk/Volume Group Space**

Amount of consumed disk or volume group space.

**Consumed File System or Logical Volume Space**

Amount of space on the file system or logical volume consumed by the computer on a disk. For computer groups, this is the amount of space on file systems or logical volumes that is consumed by the computers within those groups. For file system groups, this is the amount of space on file systems or logical volumes within those groups that is consumed. For disk groups and volume groups, this is the external logical volume or file system space consumed within those groups.

**Correlated**

Indicates **Yes** if Tivoli Storage Productivity Center is able to relate the host storage directly to a storage subsystem. You can relate this correlated storage back to the actual storage subsystem storage for a more detailed view of how the host storage is allocated.

If a computer disk, file system, or logical volume appears in the Disk Storage On Storage Subsystem reports but is not correlated, it might be the result of the following reasons:

- The storage subsystem was not probed. Run a probe to collect data about the storage subsystem then generate this report.
- The storage subsystem storage was unassigned from the host storage, but the host has not been probed. Run a probe to collect data about the host then generate this report.

- The storage subsystem storage was not correctly identified by a probe of the host. If this is the problem, check the probe log of the host for messages relating to SCSI command errors and the probe results of the storage subsystem.

**Creator**

ID of the user who created a group.

**Disk Path**

Path name of a disk as seen by the computer.

**Disk Space**

Total amount of space allocated to a computer disk. For computer groups, this is the total amount of space allocated to the disks on the computers within those groups. For disk and volume groups, this is the external disk space within the disk group or volume group.

**Note:**

- This space is distributed across the hypervisor disk space from where the virtual machine disk's file resides. The portion of space identified to a particular external hypervisor disk is reported only. This is true for all space values in this report.
- The XIV system GUI and the XIV system command-line interface (CLI) report sizes in GB of disks, volumes, pools, and so on, based on  $10^9$ . Tivoli Storage Productivity Center reports sizes in GB based on  $2^{30}$ . For example, a volume showing a size of 17 GB in the XIV system GUI is reported as 16 GB ( $\text{XIV GB size} * 1000 * 1000 * 1000 / 1024 / 1024 / 1024$  rounded to the nearest GB) in Tivoli Storage Productivity Center.

**Disk/Volume Group Name**

Name of a disk group or volume group where a logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

**Disk/Volume Group Space**

Total amount of disk group or volume group space. This is space inclusive of all storage space and applies to all space values related to disk groups or volume groups.

**Disk/Volume Group Type**

Type of a disk group.

**Fibre Attached**

Indicates **Yes** when Tivoli Storage Productivity Center determines that a disk is behind a fibre channel adapter port.

**File System Type**

Type of a file system, such as NTFS.

**File System/Logical Volume Path**

Name of a file system or logical volume path.

**Group** Name of a group.

**Grown Defects**

Number of defects detected on the disk since it was new.

**Hypervisor Name**

Name of the hypervisor where a virtual machine was created. This column displays **N/A** if not applicable to a disk.

**Logical Volume Space**

Total amount of space allocated to the logical volume on a disk. For computer groups, this is the total amount of space allocated to the logical volumes on disks within those groups. For file system groups, this is the total amount of space allocated to the logical volumes within those groups. For disk groups and volume groups, this is the external logical volume space within those groups.

This space includes the file system space on the disk reported if the entity displayed is a file system. The actual size of the file system on the disk reported is in addition to the consumed and available space.

**Note:** When relating a monitored entity to distributed file system/logical volume space, the space is distributed across the disks where a logical volume or file system resides. The portion identified to a disk is reported only.

**(Disk) Manufacturer**

The manufacturer of a disk.

**Mapped to Storage Subsystem**

Indicates **Yes** if Tivoli Storage Productivity Center has identified a disk as being located on a storage volume from a storage subsystem. The storage subsystem might not be monitored by Tivoli Storage Productivity Center.

**(Disk) Model**

Model number of a disk.

**Number of Disks**

Number of disks in a computer group. For disk groups and volume groups, this is the number of external disks within those groups.

**Number of File Systems/Logical Volumes**

Number of logical file systems or logical volumes within a group. For disk groups and volume groups, this is the number of external logical file systems or logical volumes within those groups.

The logical volumes underlying the file systems are not counted for the value in the column.

**Overallocated**

Indicates **Yes** if Tivoli Storage Productivity Center is unable to uniquely identify disk storage across more than one host. Check the probe log of the host for messages relating to SCSI command errors.

**Path Type**

Path type of the entity being reported: file system or a logical volume.

**Primary Defects**

Number of defects on the disk when it was new.

**Shared**

Indicates **Yes** if a file system or logical volume is shared by more than one computer.

**Serial Number**

Serial number of a disk.

**Unavailable Disk Space**

Amount of space on the disk not available to be consumed. For the computer groups, this is the amount of space on disks that is not available

to be consumed by the computers in a group. For disk and volume groups, this is the external disk space that is not available for use within those groups.

This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

#### **Unavailable Logical Volume Space**

Amount of space on a logical volume that is not available for use by a computer. For computer groups, this is the amount of space on logical volumes that is not available for use by the computers within those groups. For file system groups, this is the amount of space on the logical volumes within those groups that is not available for use. For disk groups and volume groups, this is the external logical volume space that is not available for use within those groups.

This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

#### **VM Disk File**

Virtual machine disk's file on a monitored hypervisor that represents the disk on the virtual machine. This column displays **N/A** if not applicable to a disk.

#### **VMFS Disk**

Hypervisor disk where the virtual machine disk resides. This column displays **N/A** if not applicable to a disk.

**Note:** If a virtual machine disk's file spans multiple hypervisor disks, duplicate entries for the virtual machine disk are generated. If the report is distributed, the space is distributed across the hypervisor disks that are not external.

#### **VFMS Mount Point**

Hypervisor file system's mount point where the virtual machine disk's file resides. This column displays **N/A** if not applicable to a disk.

#### **Related Topics**

- All Disk Storage reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

*Disk Storage - Disk Storage on Storage Subsystems reports:*

Use these reports to view information about computer disks that IBM Tivoli Storage Productivity Center has identified as storage subsystem volumes.

See Disks On Storage Subsystems - Column Descriptions for descriptions of all the columns that appear in Disk Storage on Storage Subsystem reports.

Use **By** clauses to filter the entities that appear in a report. For example, if you want to view a report that displays specific computers, select **By Storage Subsystem** and use the **Selection** and **Filter** options to determine which storage subsystems appear in the generated report.

You can customize the information in each By clause by relating the monitored entities to other storage. The following options for relating storage entities (storage subsystems, computers, groups, file systems, logical volumes, disk groups, volume groups) to other storage in these reports includes the following:

- Use the **Relate Storage Subsystems To:** options in the By Storage Subsystem and By Storage Subsystem Group reports to relate storage subsystems to the following storage:
  - Computer Space
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space
  - Computer Disk Space
  - Distributed Computer Disk Space
- Use the **Relate Computers To:** and **Relate Disk/Volume Group To:** options in the By Computer, By Computer Group, and By Disk/Volume Group reports to relate storage computers, computer groups, disk groups, or volume groups to the following storage:
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space
  - Computer Disk Space
  - Distributed Computer Disk Space
  - Storage Subsystem Space (not available for disk groups and volume groups)
- Use the **Relate FileSystems/Logical Volumes To:** options in the By File System/Logical Volume and By Filesystem Group reports to relate file systems or logical volumes to the following storage:
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space

### Considerations

Keep in mind the following considerations when generating Disk Storage on Storage Subsystems reports:

#### **By Storage Subsystem (Relate Storage Subsystems To: File System/Logical Volume Space, Distributed File System/Logical Volume Space)**

Duplicate entries appear in By Storage Subsystem reports for file systems or logical volumes that span multiple storage subsystems. The storage space shown in the report is adjusted in the totals. The storage space values include space that is not storage subsystem storage if a file system or logical volume spans disks that are both storage subsystem disks and other storage disks.

When relating storage subsystems or computers to distributed file system/logical volume space, the reports displays duplicate entries for file systems or logical volumes that span multiple disks, but the displayed storage space unique. The values for storage space exclude space that is not storage subsystem storage if a file system or logical volume spans disks that are both storage subsystem disks and other storage disks.

#### **By Storage Subsystem (Relate Storage Subsystems To: Computer Disk Space, Distributed Computer Disk Space)**

Duplicate entries appear in these reports for virtual machines whose disks spans multiple hypervisor disks. The storage space shown in the report is adjusted in the totals. The storage space values include space that is not

storage subsystem storage if the file system on a virtual machine spans hypervisor disks that are external or if the file system spans other storage disks.

When relating storage subsystems or computers to distributed computer disk space, this report displays duplicate entries for virtual machines whose disks spans multiple hypervisor disks, but the storage space shown in the report is unique. The storage space values exclude space that is not storage subsystem storage if the virtual machine disk's file system spans hypervisor disks that are external or the virtual machine disk's file system spans other storage disks.

#### **By XIV system**

The XIV system GUI and the XIV system command-line interface (CLI) report sizes in GB of disks, volumes, pools, and so on, based on  $10^9$ . Tivoli Storage Productivity Center reports sizes in GB based on  $2^{30}$ . For example, a volume showing a size of 17 GB in the XIV system GUI is reported as 16 GB (XIV GB size \* 1000 \* 1000 \* 1000 / 1024 / 1024 / 1024 rounded to the nearest GB) in Tivoli Storage Productivity Center.

#### **Related Topics**

Keep in mind the following considerations when generating Disk Storage on Storage Subsystems reports:

- Disks On Storage Subsystems - Column Descriptions
- Monitored Computer Storage Space reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

*Disks On Storage Subsystems - Column Descriptions:*

Use these reports to view information about disks on monitored computers that IBM Tivoli Storage Productivity Center has identified as storage subsystem volumes.

The following buttons and columns appear in the By clauses of Disks On Storage Subsystems reports:

**Note:** Not all of the following columns appear in each By clause.

#### **Selection...**

Click this button to select the specific entities that you want to include in the report. For example, in the **By Storage Subsystem** report, click this button to select the specific storage subsystems upon which you want to report.

**Filter** Click this button to apply conditions to the columns within a report to further select or restrict the entities that you want to include in the report.

#### **Relate <entity> To:**

Select how you want to relate monitored storage entities in the report to other storage. See Disk Storage - Disk Storage Not On Storage Subsystems reports for more information.

### Magnifying glass icon

This button appears on all **By <entity> Group** reports. Click this button to view a more detailed report about the entities in a group.

### Available Disk Space

Amount of disk space on a storage subsystem that is available for a computer to consume. For computer groups, this is the amount of space on disks that is available for use by the computers within those groups. For storage subsystem groups, this is the amount of space on disks that is available for use by the storage subsystems within those groups. For disk and volume groups, this is the external disk space available for use within those groups.

This space includes the file system space that is available on the disk and the space that has not been allocated to a logical volume.

### Available Disk/Volume Group Space

Amount of available space in a disk group or volume group.

### Available File System/Logical Volume Space

Amount of space on the file system or logical volume that is available for use by the computer on a storage subsystem. For computers groups and storage subsystem groups, this is the amount of space on file systems or logical volumes that is available for use on storage subsystems by the entities within those groups. For file system groups, this is the amount of space on file systems or logical volumes within those groups that can be consumed. For disk groups and volume groups, this is the external logical volume or file system space available for use within those groups.

### Computer

Name of a computer where a disk, file system, logical volume, or disk group is located.

### Computer OS Type

The operating system running on a computer.

### Consumed Disk Space

Amount of disk space on a storage subsystem that is consumed by a computer. For storage subsystem groups, this is the amount of consumed disk space within those groups. For computer groups, this is the amount of disk space consumed by the computers those groups. For disk and volume groups, this is the external disk space consumed within those groups.

This space includes the file system space that is consumed on a disk. If a logical volume on this disk space has no file system, then this space includes the consumed portion of the logical volume that has been allocated to the disk.

### Consumed Disk/Volume Group Space

Amount of consumed disk or volume group space.

### Consumed File System/Logical Volume Space

Amount of space on the file system or logical volume that is consumed by the computer on a storage subsystem. For computer groups and storage subsystem groups, this is the amount of space on file systems or logical volumes that is consumed by the entities within those groups. For file system groups, this is the amount of storage subsystem space on file systems or logical volumes within those groups that is consumed. For disk groups and volume groups, this is the external logical volume or file system space consumed within those groups.



**Correlated**

Indicates **Yes** if Tivoli Storage Productivity Center is able to relate the host storage directly to a storage subsystem. You can relate this correlated storage back to the actual storage subsystem storage for a more detailed view of how the host storage is allocated.

If a computer disk, file system, or logical volume appears in the Disk Storage On Storage Subsystem reports but is not correlated, it might be the result of the following reasons:

- The storage subsystem was not probed. Run a probe to collect data about the storage subsystem then generate this report.
- The storage subsystem storage was unassigned from the host storage, but the host has not been probed. Run a probe to collect data about the host then generate this report.
- The storage subsystem storage was not correctly identified by a probe of the host. If this is the problem, check the probe log of the host for messages relating to SCSI command errors and the probe results of the storage subsystem.

**Creator**

ID of the user who created a group.

**Disk Path**

Path name of a disk as seen by the computer.

**Disk Space**

Total disk space on a storage subsystem seen by a computer. For computer groups, this is the total disk space within those groups that is associated with storage subsystems. For storage subsystem groups, this is the total disk space on the storage subsystems within those groups. For disk and volume groups, this is the external disk space within the disk group or volume group.

**Note:** This space is distributed across the hypervisor disk space from where the virtual machine disk's file resides. The portion of space identified to a particular external hypervisor disk is reported only. This is true for all space values in this report.

**Disk/Volume Group Name**

Name of a disk group or volume group where a logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

**Disk/Volume Group Space**

Total amount of disk group or volume group space. This is space inclusive of all storage space and applies to all space values related to disk groups or volume groups.

**Disk/Volume Group Type**

Type of a disk group. For example, a volume group might be LVM or Veritas. This column displays **N/A** if a disk group or volume group does not exist on a disk.

**Fibre Attached**

Indicates **Yes** when Tivoli Storage Productivity Center determines that a disk is behind a fibre channel adapter port.

**File System Type**

Type of a file system, such as NTFS.

**File System/Logical Volume Path**

Name of a file system or logical volume path.

**Group** Name of a group.

**Grown Defects**

Number of defects detected on the disk since it was new.

**Hypervisor Name**

Name of the hypervisor where a virtual machine was created. This column displays **N/A** if not applicable to a disk.

**Logical Volume Space**

Total amount of space allocated to the logical volume on a disk. For computer groups, this is the total logical volume space in computer groups that is associated with storage subsystems. For storage subsystem groups, this is the total logical volume space on the storage subsystems within those groups. For file system groups, this is the total amount of space allocated to the logical volumes within those groups. For disk groups and volume groups, this is the external logical volume space within those groups.

This space includes the file system space on the disk reported if the entity displayed is a file system. The actual size of the file system on the disk reported is in addition to the consumed and available space.

**Note:** When relating a monitored entity to distributed file system/logical volume space, the space is distributed across the disks where a logical volume or file system resides. The portion identified to a disk is reported only.

**(Disk) Manufacturer**

The manufacturer of a disk.

**Mapped to Storage Subsystem**

Indicates **Yes** if Tivoli Storage Productivity Center has identified a disk as being located on a storage volume from a storage subsystem. The storage subsystem might not be monitored by Tivoli Storage Productivity Center.

**(Disk) Model**

Model number of a disk.

**Number of Disks**

Number of disks assigned to a storage subsystem or group. For disk groups and volume groups, this is the number of external disks within those groups.

**Number of File Systems/Logical Volumes**

Number of logical file systems or logical volumes within a group. For disk groups and volume groups, this is the number of external logical file systems or logical volumes within those groups.

The logical volumes underlying the file systems are not counted for the value in the column.

**Overallocated**

Indicates **Yes** if Tivoli Storage Productivity Center is unable to uniquely identify disk storage across more than one host. Check the probe log of the host for messages relating to SCSI command errors.

**Path Type**

Path type of the entity being reported: file system or a logical volume.

**Primary Defects**

Number of defects on the disk when it was new.

**Shared**

Indicates **Yes** if a file system or logical volume is shared by more than one computer.

**Serial Number**

Serial number of a disk.

**Storage Subsystem Manufacturer**

The manufacturer of a storage subsystem.

**Storage Subsystem Type**

The type of storage subsystem.

**Unavailable Disk Space**

Amount of disk space on the storage subsystem that is not available for a computer to consume. For the computer groups, this is the amount of space on disks that is not available to be consumed by the computers within those groups. For the storage subsystem groups, this is the amount of space on disks that is not available to be consumed within those groups. For disk and volume groups, this is the external disk space that is not available for use within those groups.

This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

**Unavailable Logical Volume Space**

Amount of space on the logical volume that is not available for use by the computer on a storage subsystem. For computer groups and storage subsystem groups, this is the amount of space on logical volumes that is not available for use by the entities within those groups. For file system groups, this is the amount of space on the logical volumes within those groups that is not available for use. For disk groups and volume groups, this is the external logical volume space that is not available for use within those groups.

This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

**VM Disk File**

Virtual machine disk's file on a monitored hypervisor that represents the disk on the virtual machine. This column displays **N/A** if not applicable to a disk.

**VMFS Disk**

Hypervisor disk where the virtual machine disk resides. This column displays **N/A** if not applicable to a disk.

**Note:** If a virtual machine disk's file spans multiple hypervisor disks, duplicate entries for the virtual machine disk are generated. If the report is distributed, the space is distributed across the hypervisor disks that are not external.

**VFMS Mount Point**

Hypervisor file system's mount point where the virtual machine disk's file resides. This column displays **N/A** if not applicable to a disk.

**Related Topics**

- Disks On Storage Subsystems reports overview

- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

*Disk Storage - Disk Storage Not On Storage Subsystems reports:*

Use these reports to view information about computer disks that IBM Tivoli Storage Productivity Center has not identified as storage subsystem volumes.

These reports are organized into the following categories:

#### **Disks With Serial Numbers**

These reports display information about non-storage subsystem disks that Tivoli Storage Productivity Center has identified with a valid serial number.

#### **Disks Without Serial Numbers**

These reports display information about non-storage subsystem disks that Tivoli Storage Productivity Center has not identified with serial numbers that uniquely identify those disks. This might occur on internal disks within a hypervisor and virtual machines that have storage created from internal disks within a hypervisor. If this problem occurs for other types of storage, check the probe log of the host for messages relating to SCSI command errors.

See Disks Not On Storage Subsystems - Column Descriptions for descriptions of all the columns that appear in Disk Storage Not On Storage Subsystems reports.

Use By clauses to filter the entities that appear in a report. For example, if you want to view a Disk Storage Not On Storage Subsystems report that displays specific computers, select **By Computer** and use the **Selection** and **Filter** options to determine which computers appear in that report.

You can customize the information in each By clause by relating the monitored entities to other storage. The following options for relating storage entities (computers, computer groups, file systems, file system groups, logical volumes, disk groups, volume groups) to other storage in these reports includes the following:

- Use the **Relate Computers To:** and **Relate Disk/Volume Group To:** options in the By Computer, By Computer Group, and By Disk/Volume Group reports to relate computers, computer groups, disk groups, and volume groups to the following storage:
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space
  - Computer Disk Space
  - Distributed Computer Disk Space
- Use the **Relate FileSystems/Logical Volumes To:** options in the By File System/Logical Volume and By Filesystem Group reports to relate file systems or logical volumes to the following storage:
  - File System or Logical Volume Space
  - Distributed File System or Logical Volume Space

## Related Topics

- Disks Not On Storage Subsystems - Column Descriptions
- Monitored Computer Storage Space reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

### *Disks Not On Storage Subsystems - Column Descriptions:*

Use Disks Not On Storage Subsystems reports to view information about disk storage that is not associated with storage subsystems. IBM Tivoli Storage Productivity Center has not identified the disks in this report as storage subsystem volumes. Use this topic to view descriptions of the columns that appear in these reports.

The following buttons and columns appear in the By clauses of Disks Not On Storage Subsystems reports:

**Note:** Not all of the following columns appear in each By clause.

#### **Selection...**

Click this button to select the specific entities that you want to include in the report. For example, in the **By Filesystem/Logical Volume** report, click this button to select the specific file systems or logical volumes upon which you want to report.

**Filter** Click this button to apply conditions to the columns within a report to further select or restrict the entities that you want to include in the report.

#### **Relate <entity> To:**

Select how you want to relate monitored storage entities in the report to other storage. See Disk Storage - Disk Storage Not On Storage Subsystems reports for more information.

#### **Magnifying glass icon**

This button appears on all **By <entity> Group** reports. Click this button to view a more detailed report about the entities in a group.

#### **Available Disk Space**

Amount of space on a disk that is available for use by a computer. For computer groups, this is the amount of space on disks that is available for use by the computers in a group. For disk and volume groups, this is the external disk space available for use within those groups.

This space includes the file system space that is available on the disk and the space that has not been allocated to a logical volume.

#### **Available Disk/Volume Group Space**

Amount of available space in a disk group or volume group.

#### **Available File System/Logical Volume Space**

Amount of space on the file system or logical volume that is available for use on a disk. For computers groups, this is the amount of space on file systems or logical volumes that is available for use by the computers within those groups. For file system groups, this is the amount of space on file systems or logical volumes within those groups that can be consumed.

For disk groups and volume groups, this is the external logical volume or file system space available for use within those groups.

**Computer**

Name of a computer where a disk, file system, logical volume, or disk group is located.

**Computer OS Type**

The operating system running on a computer.

**Consumed Disk Space**

Amount of space on a disk that is consumed. For computer groups, this is the amount of disk consumed by the computers in those groups. For disk and volume groups, this is the external disk space consumed within those groups.

This space includes the file system space that is consumed on a disk. If a logical volume on this disk space has no file system, then this space includes the consumed portion of the logical volume that has been allocated to the disk.

**Consumed Disk/Volume Group Space**

Amount of consumed disk or volume group space.

**Consumed File System/Logical Volume Space**

Amount of space on the file system or logical volume consumed by the computer on a disk. For computer groups, this is the amount of space on file systems or logical volumes that is consumed by the computers within those groups. For file system groups, this is the amount of space on file systems or logical volumes within those groups that is consumed. For disk groups and volume groups, this is the external logical volume or file system space consumed within those groups.

**Creator**

ID of the user who created a group.

**Disk Path**

Path name of a disk as seen by the computer.

**Disk Space**

Total amount of space allocated to a computer disk. For computer groups, this is the total amount of space allocated to the disks on the computers within those groups. For disk and volume groups, this is the external disk space within the disk group or volume group.

**Note:** This space is distributed across the hypervisor disk space from where the virtual machine disk's file resides. The portion of space identified to a particular external hypervisor disk is reported only. This is true for all space values in this report.

**Disk/Volume Group Name**

Name of a disk group or volume group where a logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

**Disk/Volume Group Space**

Total amount of disk group or volume group space. This is space inclusive of all storage space and applies to all space values related to disk groups or volume groups.

**Disk/Volume Group Type**

Type of a disk group. For example, a volume group might be LVM or Veritas. This column displays **N/A** if a disk group or volume group does not exist on a disk.

**Fibre Attached**

Indicates **Yes** when Tivoli Storage Productivity Center determines that a disk is behind a fibre channel adapter port.

**File System Type**

Type of a file system, such as NTFS.

**File System/Logical Volume Path**

Name of a file system or logical volume path.

**Group** Name of a group.**Grown Defects**

Number of defects detected on the disk since it was new.

**Hypervisor Name**

Name of the hypervisor where a virtual machine was created. This column displays **N/A** if not applicable to a disk.

**Logical Volume Space**

Total amount of space allocated to the logical volume on a disk. For computer groups, this is the total amount of space allocated to the logical volumes on disks within those groups. For file system groups, this is the total amount of space allocated to the logical volumes within those groups. For disk groups and volume groups, this is the external logical volume space within those groups.

This space includes the file system space on the disk reported if the entity displayed is a file system. The actual size of the file system on the disk reported is in addition to the consumed and available space.

**Note:** When relating a monitored entity to distributed file system/logical volume space, the space is distributed across the disks where a logical volume or file system resides. The portion identified to a disk is reported only.

**(Disk) Manufacturer**

The manufacturer of a disk.

**(Disk) Model**

Model number of a disk.

**Number of Disks**

Number of disks in a computer group. For disk groups and volume groups, this is the number of external disks within those groups.

**Number of File Systems/Logical Volumes**

Number of logical file systems or logical volumes within a group. For disk groups and volume groups, this is the number of external logical file systems or logical volumes within those groups.

The logical volumes underlying the file systems are not counted for the value in the column.

**Overallocated**

Indicates **Yes** if Tivoli Storage Productivity Center is unable to uniquely identify disk storage across more than one host. Check the probe log of the host for messages relating to SCSI command errors.

**Path Type**

Path type of the entity being reported: file system or a logical volume.

**Primary Defects**

Number of defects on the disk when it was new.

**Shared**

Indicates **Yes** if a file system or logical volume is shared by more than one computer.

**Serial Number**

Serial number of a disk.

**Unavailable Disk Space**

Amount of space on the disk not available to be consumed. For the computer groups, this is the amount of space on disks that is not available to be consumed by the computers in a group. For disk and volume groups, this is the external disk space that is not available for use within those groups.

This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

**Unavailable Logical Volume Space**

Amount of space on a logical volume that is not available for use by a computer. For computer groups, this is the amount of space on logical volumes that is not available for use by the computers within those groups. For file system groups, this is the amount of space on the logical volumes within those groups that is not available for use. For disk groups and volume groups, this is the external logical volume space that is not available for use within those groups.

This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

**VM Disk File**

Virtual machine disk's file on a monitored hypervisor that represents the disk on the virtual machine. This column displays **N/A** if not applicable to a disk.

**VMFS Disk**

Hypervisor disk where the virtual machine disk resides. This column displays **N/A** if not applicable to a disk.

**Note:** If a virtual machine disk's file spans multiple hypervisor disks, duplicate entries for the virtual machine disk are generated. If the report is distributed, the space is distributed across the hypervisor disks that are not external.

**VFMS Mount Point**

Hypervisor file system's mount point where the virtual machine disk's file resides. This column displays **N/A** if not applicable to a disk.

**Related Topics**

- Disks Not On Storage Subsystems reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.



## Disk Storage - Non-Disk Storage reports:

Use these reports to view information about storage space that does not reside on disks that are monitored by IBM Tivoli Storage Productivity Center. This includes storage that resides on a remote mount or a network attached server and storage that is temporarily allocated to processor memory.

Use **By** clauses to filter the entities that appear in a report. For example, if you want to view a report that displays specific computers, select **By Computer** and use the **Selection** and **Filter** options to determine which computers appear in the generated report.

The following **By** clauses are available for Non-Disk Storage reports:

- By Computer
- By Computer Group
- By Filesystem/Logical Volume
- By Filesystem Group
- By Disk/Volume Group

You can customize the information in each **By** clause by relating the monitored entities to other storage. The following options for relating storage entities (computers, computer groups, file systems, file system groups, logical volumes, disk groups, volume groups) to other storage in these reports includes the following:

- File System or Logical Volume Space
- Distributed File System or Logical Volume Space

### Related Topics

- Monitored Computer Storage Space reports overview
- See Storage entities to view descriptions of the storage entities that appear in these reports.
- See Common storage space values to view descriptions of the common storage space values that appear in these reports.

*Non-Disk Storage - By Computer report, By Filesystem/Logical Volume report:*

Use these reports to view information about non-disk storage. This includes storage that resides on a remote mount or a network attached server and storage that is temporarily allocated to processor memory. Select **By Computer** if you want to determine which computers are displayed in the report. Select **By Filesystem/Logical Volume** if you want to determine which file systems or logical volumes appear in the report.

### Selection...

For the **By Computer** reports, click this button to select the specific computers upon which you want to report. For the **By Filesystem/Logical Volume** reports, click this button to select the specific file systems or logical volumes upon which you want to report.

**Filter** Click this button to apply conditions to the columns within a report to further select or restrict computers or file system/logical volumes upon which you want to report.

**Relate <entity> To:**

Select how you want to relate computers or file systems in the report to other storage entities. See JDisk Storage - Non-Disk Storage reports for more information.

**Computer**

Name of a computer.

**File System/Logical Volume Path**

Name of a file system or logical volume path.

**Note:** Each monitored file system or logical volume path is displayed only once in this report because if the storage space is not distributed (there are no underlying disks), then there is no need to report distribution.

**Logical Volume Space**

Total amount of space allocated to the logical volume on a disk. This space includes the file system space if the entity displayed is a file system. The actual size of the file system is in addition to the consumed and available space.

**Note:** When relating computers to distributed file system/logical volume space, the space is distributed across the logical volumes where the file system or parent logical volume resides. The portion identified to a particular file system or parent logical volume is reported only. If a file system does not have an underlying logical volume, the logical volume space is the size of that file system.

**Consumed File System/Logical Volume Space**

Amount of space on a file system or logical volume that is consumed by a computer.

**Available File System/Logical Volume Space**

Amount of space on a file system or logical volume that is available for use by a computer.

**Unavailable Logical Volume Space**

Amount of space on the logical volume that is not available for use by a computer. This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

**Computer OS Type**

Operating system running on a computer.

**Disk/Volume Group Name**

Name of a disk or volume group where the logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

**Disk/Volume Group Type**

Type of a disk or volume group where the logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

**File System Type**

Type of a file system, such as NTFS.

**Path Type**

Path type of an entity: file system or a logical volume.

**Shared**

Indicates **Yes** if a file system or logical volume is shared by more than one computer.

*Non-Disk Storage - By Computer Group report, By Filesystem Group report:*

Use these reports to view information about non-disk storage. This includes storage that resides on a remote mount or a network attached server and storage that is temporarily allocated to processor memory. Select **By Computer Group** if you want to determine which computer groups are displayed in the report. Select **By Filesystem Group** if you want to determine which file system groups appear in the report.

**Selection...**

For the **By Computer Group** reports, click this button to select the specific computer groups upon which you want to report. For the **By Filesystem Group** reports, click this button to select the specific file system group upon which you want to report.

**Filter** Click this button to apply conditions to the columns within a report to further select or restrict groups upon which you want to report.

**Relate <entity> To:**

Select how you want to relate groups in the report to other storage entities. See Disk Storage - Non-Disk Storage reports for more information.

A unique row appears for each group in this report.

**Magnifying glass icon** 

Click this button to view a more detailed report about the computers or file systems in a group.

**Creator**

ID of the user who created the group.

**Group** Name of a group.

**Logical Volume Space**

Total amount of space allocated to the computers or file systems in a group. This space includes the file system space if the entity displayed is a file system. The actual size of the file system is in addition to the consumed and available space.

**Note:** When relating groups to distributed file system/logical volume space, this space is distributed across the disks from where the logical volume or file system reside. This is true for all space values in the Distributed File System/Logical Volume Space report.

**Consumed File System/Logical Volume Space**

Amount of space on file systems or logical volumes within a group that is consumed.

**Available File System/Logical Volume Space**

Amount of space on file systems or logical volumes within a group that can be consumed.

**Unavailable Logical Volume Space**

Amount of space on the logical volumes in a group that is not available for use. This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

## Number of File Systems/Logical Volumes

Number of file systems or logical volumes in a group.

*Non-Disk Storage - By Computer report, By Filesystem/Logical Volume report:*

Use these reports to view information about non-disk storage. This includes storage that resides on a remote mount or a network attached server and storage that is temporarily allocated to processor memory. Select **By Computer** if you want to determine which computers are displayed in the report. Select **By Filesystem/Logical Volume** if you want to determine which file systems or logical volumes appear in the report.

### Selection...

For the **By Computer** reports, click this button to select the specific computers upon which you want to report. For the **By Filesystem/Logical Volume** reports, click this button to select the specific file systems or logical volumes upon which you want to report.

**Filter** Click this button to apply conditions to the columns within a report to further select or restrict computers or file system/logical volumes upon which you want to report.

### Relate <entity> To:

Select how you want to relate computers or file systems in the report to other storage entities. See JDisk Storage - Non-Disk Storage reports for more information.

### Computer

Name of a computer.

### File System/Logical Volume Path

Name of a file system or logical volume path.

**Note:** Each monitored file system or logical volume path is displayed only once in this report because if the storage space is not distributed (there are no underlying disks), then there is no need to report distribution.

### Logical Volume Space

Total amount of space allocated to the logical volume on a disk. This space includes the file system space if the entity displayed is a file system. The actual size of the file system is in addition to the consumed and available space.

**Note:** When relating computers to distributed file system/logical volume space, the space is distributed across the logical volumes where the file system or parent logical volume resides. The portion identified to a particular file system or parent logical volume is reported only. If a file system does not have an underlying logical volume, the logical volume space is the size of that file system.

### Consumed File System/Logical Volume Space

Amount of space on a file system or logical volume that is consumed by a computer.

### Available File System/Logical Volume Space

Amount of space on a file system or logical volume that is available for use by a computer.

### Unavailable Logical Volume Space

Amount of space on the logical volume that is not available for use by a

computer. This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

**Computer OS Type**

Operating system running on a computer.

**Disk/Volume Group Name**

Name of a disk or volume group where the logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

**Disk/Volume Group Type**

Type of a disk or volume group where the logical volume or file system was created. This column displays **N/A** if a disk group or volume group does not exist on a disk.

**File System Type**

Type of a file system, such as NTFS.

**Path Type**

Path type of an entity: file system or a logical volume.

**Shared**

Indicates **Yes** if a file system or logical volume is shared by more than one computer.

*Non-Disk Storage - By Computer Group report, By Filesystem Group report:*

Use these reports to view information about non-disk storage. This includes storage that resides on a remote mount or a network attached server and storage that is temporarily allocated to processor memory. Select **By Computer Group** if you want to determine which computer groups are displayed in the report. Select **By Filesystem Group** if you want to determine which file system groups appear in the report.

**Selection...**

For the **By Computer Group** reports, click this button to select the specific computer groups upon which you want to report. For the **By Filesystem Group** reports, click this button to select the specific file system group upon which you want to report.

**Filter** Click this button to apply conditions to the columns within a report to further select or restrict groups upon which you want to report.

**Relate <entity> To:**

Select how you want to relate groups in the report to other storage entities. See Disk Storage - Non-Disk Storage reports for more information.

A unique row appears for each group in this report.

**Magnifying glass icon** 

Click this button to view a more detailed report about the computers or file systems in a group.

**Creator**

ID of the user who created the group.

**Group** Name of a group.

**Logical Volume Space**

Total amount of space allocated to the computers or file systems in a

group. This space includes the file system space if the entity displayed is a file system. The actual size of the file system is in addition to the consumed and available space.

**Note:** When relating groups to distributed file system/logical volume space, this space is distributed across the disks from where the logical volume or file system reside. This is true for all space values in the Distributed File System/Logical Volume Space report.

**Consumed File System/Logical Volume Space**

Amount of space on file systems or logical volumes within a group that is consumed.

**Available File System/Logical Volume Space**

Amount of space on file systems or logical volumes within a group that can be consumed.

**Unavailable Logical Volume Space**

Amount of space on the logical volumes in a group that is not available for use. This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

**Number of File Systems/Logical Volumes**

Number of file systems or logical volumes in a group.

*Disk Storage on Storage Subsystems reports - By Disk/Volume Group:*

Use this report to view information about non-disk storage. This includes storage that resides on a remote mount or a network attached server and storage that is temporarily allocated to processor memory. This report enables you to determine which disk groups or volume groups are displayed.

**Selection...**

Click this button to select the specific disk groups or volume groups upon which you want to report.

**Filter**

Click this button to apply conditions to the columns within a report to further select or restrict disk groups or volume groups upon which you want to report.

**Relate Disk/Volume Group To:**

Select how you want to relate the non-disk storage in disk groups or volume groups to other storage entities. See Disk Storage - Non-Disk Storage reports for more information.

A unique row appears for each disk group or volume group.

**Magnifying glass icon** 

Click this button to view a more detailed report about the disk or file system within the disk group or volume group.

**Disk/Volume Group Name**

Name of a group.

**Disk/Volume Group Type**

Type of a group. For example, a volume group might be LVM or Veritas.

**Computer**

Name of the computer where a disk group or volume group resides.

**Computer OS Type**

Operating system running on a computer.

**Disk/Volume Group Space**

Total amount of disk group or volume group space. This is space inclusive of all storage space and applies to all space values related to disk groups or volume groups.

**Consumed Disk/Volume Group Space**

Amount of consumed disk group or volume group space.

**Available Disk/Volume Group Space**

Amount of available disk group or volume group space.

**Logical Volume Space**

External logical volume space within a disk group or volume group.

**Consumed File System/Logical Volume Space**

External logical volume or file system space that is consumed within a disk group or volume group.

**Available File System/Logical Volume Space**

External logical volume or file system space that is available for use within a disk group or volume group.

**Unavailable Logical Volume Space**

External logical volume space that is not available for use within a disk group or volume group. This space is used for logical volume overhead and is reserved by the operating system to manage a file system.

**Number of File Systems/Logical Volumes**

Number of external logical file systems or logical volumes within a disk or volume group. The logical volumes underlying the file systems are not counted.

**Additional reporting support and information**

Learn about the additional storage devices upon which you can report within Data Manager.

**IBM FAStT Subsystem support:** The IBM FAStT subsystem reporting feature further extends the subsystem reporting capabilities of Data Manager. You can use the Reporting facility to view information about the capacity, allocation, and volumes of an IBM FAStT subsystem. You can view IBM FAStT subsystem information by using the following menu options to generate reports:

- **Reporting → Asset → By Storage Subsystem → <IBM FAStT subsystem name>**
- **Reporting → Asset → System-wide → Storage Subsystems**
- **Reporting → Asset → System-wide → Volumes**
- **Reporting → Asset → System-wide → Disks**

**Note:** Data Manager shows an incorrect value for the **Storage Pool Free Space** column in IBM FAStT subsystem reports (for example, **Asset → By Storage Subsystem → <FAStT\_subsystem\_name>**, and **Asset → System-wide → Storage Subsystems**). The value for **Storage Pool Free Space** in these reports reflects the space from one Free Capacity node rather than all the free space within the defined arrays on a FAStT subsystem.

**Hitachi Data Systems Subsystem support:** Use the Reporting facility to report on Hitachi Data Systems (HDS) storage subsystems. You can view reports that contain information about the capacity, allocation, and volumes of an HDS subsystem.

Any storage pools or volumes created using the HDS **HiCommand** software will not be reported through the HDS CIMOM provider. **HiCommand** is the native interface for HDS subsystems. If you configure an HDS subsystem through **HiCommand**, IBM Tivoli Storage Productivity Center will not be able to gather information about storage pools and volumes from that subsystem. The HDS CIM agent provider will return all storage pools created through the **HiCommand** interface as a single pool containing all of the **HiCommand** created pools. This single pool is always reported as Raid 5.

**Note:** Data Manager does not report on the raw capacity of disks internal to the HDS subsystem. Information on the internal disks of the HDS subsystem is not available through the CIMOM interface provided by the Hitachi HiCommand Device Manager.

You can view HDS storage subsystem information by using the following menu options to generate reports:

- **Reporting → Asset → By Storage Subsystem → <HDS subsystem name>**
- **Reporting → Asset → System-wide → Storage Subsystems**
- **Reporting → Asset → System-wide → Volumes**
- **Reporting → Asset → System-wide → Disks**
- **Reporting → TPC-wide Storage Space → Disk Space**

**HP StorageWorks Subsystem support:** Use the Reporting facility to report on HP StorageWorks storage subsystems. You can view reports that contain information about the capacity, allocation, and volumes of an HP StorageWorks subsystem.

You can view HP StorageWorks storage subsystem information by using the following menu options to generate reports:

- **Reporting → Asset → By Storage Subsystem → <HP StorageWorks subsystem name>**
- **Reporting → Asset → System-wide → Storage Subsystems**
- **Reporting → Asset → System-wide → Volumes**
- **Reporting → Asset → System-wide → Disks**

**IBM Tivoli Storage SAN File Systems support:** You can also perform scans and probes to gather information about IBM Tivoli Storage SAN File Systems and generate reports based on that information using the Reporting facility. You can view information gathered by probes of the SAN File System by using the following menu options to generate reports:

- **Reporting → Asset → By OS Type → IBM SAN File System**
- **Reporting → Asset → By Computer**
- **Reporting → TPC-wide Storage Space → File System Space**
- **Reporting → TPC-wide Storage Space → Consumed File System Space**
- **Reporting → TPC-wide Storage Space → Available File System Space**

**Note:** Due to the complexity of the SAN File System, the values displayed for total capacity, used space and free space in the Capacity and Asset reports might not be accurate. The values displayed are the values returned by the probe or scan and they may not represent actual values for your SAN File System. Use the SAN File System admin console to obtain the actual values.

You can view information gathered by scans of the SAN File System by using the following menu options to generate reports:



- **Reporting → Usage**
- **Reporting → Usage Violations**
- **Reporting → Backup**

Refer to the *Creating Directory Groups* section in the *Monitoring* chapter for more information about how to configure SAN File System support within the Reporting facility.

**Report Array, Loop, and Adapter Pair for IBM TotalStorage Enterprise Storage Server Disk Groups support:** Disk Manager provides report information on Array, Loop, and Adapter Pairs for TotalStorage Enterprise Storage Server Disk Groups. This information is available when managing an TotalStorage Enterprise Storage Server through IBM CIM Agent (ICAT) for TotalStorage Enterprise Storage Server version 1.2.0.13. Contact IBM Customer Support to obtain IBM CIM Agent for TotalStorage Enterprise Storage Server version 1.2.0.13.

You can view Array, Loop, and Adapter Pair information by using the menu option **Reporting → Asset Reporting → By Storage Subsystem** to generate reports.

#### **Conversion of volume sizes by IBM Tivoli Storage Productivity Center:**

Tivoli Storage Productivity Center uses 1024 bytes to represent a kilobyte. However, some element managers use 1000 bytes to represent a kilobyte. When reporting volume sizes, Tivoli Storage Productivity Center converts base-1000 values to base-1024 values.

The following table shows the byte values that are associated with base-1024 and base-1000 systems:

| Unit                                                                                                                                                                                                                                                                                          | Abbreviation                        | Bytes in base-1024            | Bytes in base-1000             |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|-------------------------------|--------------------------------|
| Kilobyte                                                                                                                                                                                                                                                                                      | KB in base-1024                     | $2^{10}$ or 1024              | $10^3$ or 1000                 |
|                                                                                                                                                                                                                                                                                               | kB in base-1000 <sup>see note</sup> |                               |                                |
| Megabyte                                                                                                                                                                                                                                                                                      | MB                                  | $1024^2$ or 1,048,576         | $10^6$ or 1,000,000            |
| Gigabyte                                                                                                                                                                                                                                                                                      | GB                                  | $1024^3$ or 1,073,741,824     | $10^9$ or 1,000,000,000        |
| Terabyte                                                                                                                                                                                                                                                                                      | TB                                  | $1024^4$ or 1,099,511,627,716 | $10^{12}$ or 1,000,000,000,000 |
| <p><small>Note:</small> In SI Metric Style, the abbreviation of the prefix kilo is lowercase k. IBM distinguishes between lowercase k (meaning 1000 or <math>10^3</math>) and uppercase K (meaning 1024 or <math>2^{10}</math>). Therefore, KB means 1024 bytes, and kB means 1000 bytes.</p> |                                     |                               |                                |

You can verify the volume size reported by Tivoli Storage Productivity Center against the size reported by element managers that use a base-1000 system. For example, suppose an element manager reports the size of a volume as 62,000 MB. Use the following calculations to verify the total number of base-1024 GB as reported by Tivoli Storage Productivity Center:

1. Multiply the volume size reported by the element manager by the number of bytes in a base-1000 MB:  
 $62,000 \text{ MB} \times 1,000,000 \text{ bytes per MB} = 62,000,000,000 \text{ bytes}$
2. Divide the result by the number of bytes in a base-1024 MB:  
 $62,000,000,000 \text{ bytes} / 1,048,576 \text{ bytes per MB} = 59,127.81 \text{ MB}$
3. Divide the quotient by the number of MB in a base-1024 GB:

59,127.81 MB / 1024 MB per GB = 57.74 GB

Tivoli Storage Productivity Center reports the volume size as 57.74 GB.

**VMware capacity reports:**

You must probe both the ESX Server and the Data agent or Storage Resource agent on the virtual machines before you can generate accurate reports for disk and file system capacity.

If you have an ESX Server that has 100 GB and 60 GB is allocated to the virtual machine. The virtual machine uses 5 GB of space. Both the ESX Server (H1) and the virtual machine (VC1) have been probed. You also have a physical computer (PC1) that has been probed. The TOTAL capacity for the file system or disk capacity row includes everything - virtualized disks and virtual machines as well as non-virtualized disks and machines.

| Column heading | Capacity | Used Space<br>(calculated as<br>capacity minus free<br>space) | Free Space |
|----------------|----------|---------------------------------------------------------------|------------|
| TOTAL          | 130 GB   | 25 GB                                                         | 105 GB     |
| H1             | 100 GB   | 60 GB                                                         | 40 GB      |
| VC1            | 60 GB    | 5 GB                                                          | 55 GB      |
| PC1            | 30 GB    | 20 GB                                                         | 10 GB      |

If you have probed the PC1 (physical computer) and VC1 (virtual computer) but have not probed the ESX Server (hypervisor), the capacity is shown as follows:

| Column heading | Capacity | Used Space<br>(calculated as<br>capacity minus free<br>space, any negative<br>values are<br>represented as 0) | Free Space |
|----------------|----------|---------------------------------------------------------------------------------------------------------------|------------|
| TOTAL          | 30 GB    | 0 GB                                                                                                          | 65 GB      |
| PC1            | 30 GB    | 20 GB                                                                                                         | 10 GB      |
| VC1            | 60 GB    | 5 GB                                                                                                          | 55 GB      |

If you have probed the hypervisor (H1) and PC1 (physical computer) but have not probed the VC1 (virtual computer), the capacity is shown as follows:

| Column heading | Capacity | Used Space<br>(calculated as<br>capacity minus free<br>space) | Free Space |
|----------------|----------|---------------------------------------------------------------|------------|
| TOTAL          | 130 GB   | 80 GB                                                         | 50 GB      |
| H1             | 100 GB   | 60 GB                                                         | 40 GB      |
| PC1            | 30 GB    | 20 GB                                                         | 10 GB      |

## Data Manager for Databases reports

Use Data Manager for Databases for reports to view both overview and detailed information about your RDBMS storage resources, including instances, tablespaces, databases, devices, containers, data files, fragments, tables, control files, redo logs, archive log directories, users. These reports are constructed from the statistical information gathered by data collection jobs and accumulated in the database repository.

### Using Data Manager for Databases reports

Learn how to use Data Manager for Databases reports to view information about the instances, tablespaces, databases, devices, containers, data files, fragments, tables, control files, redo logs, archive log directories, and users within your environment.

#### Before you begin:

Learn about information you should know before generating and viewing Data Manager for Databases reports.

Before you can use reports to view detailed information about tables in an RDBMS, you must have 1) defined a table group that contains the tables about which you want to view storage information and 2) scheduled and run a scan against that table group. See “Working with groups of resources” on page 371 for information about using groups. See “Collecting storage usage and trending information (scans)” on page 343 for information about scans and how to create them.

#### Data Manager for Databases - Report Categories:

Use the report categories within Data Manager for Databases reports to collect information about RDBMS storage assets, capacity, usage, and usage violations.

The following table describes the categories of reports available in Data Manager for Databases.

| Report Category | Use these reports to do the following tasks                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | What job collects information for these reports? |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------|
| Asset Reports   | <p>View the information gathered by probes. While running probes, Data Manager for Databases itemizes the information for each type of database.</p> <ul style="list-style-type: none"><li>• Oracle: instances, table spaces, tables, data files, control files, redo logs, and archive log directories</li><li>• SQL Server: instances, databases, tables, and data files</li><li>• Sybase: instances, devices, databases, tables, and fragments.</li><li>• UDB: instances, databases, table spaces, tables, containers, and users</li></ul> | Probes                                           |

| Report Category  | Use these reports to do the following tasks                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | What job collects information for these reports? |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------|
| Capacity Reports | <p>View the storage capacity of different entities within an RDBMS. Specifically, use Capacity reports to view:</p> <ul style="list-style-type: none"> <li>• <b>Instance storage capacity within all DBMSs.</b> If you are monitoring the instances for different DBMSs, you can use the <b>All DBMSs</b> node to view a summary of the storage capacity for all those instances in a single view.</li> <li>• <b>Oracle databases</b> <ul style="list-style-type: none"> <li>– How much storage you have for instances</li> <li>– How much storage is being used for data files</li> <li>– What percentage of storage for a data file is free for expansion</li> </ul> </li> <li>• <b>SQL Server databases</b> <ul style="list-style-type: none"> <li>– How much storage you have for instances</li> <li>– How much storage is being used for data files</li> <li>– What percentage of storage for a data file is free for expansion</li> </ul> </li> <li>• <b>Sybase databases</b> <ul style="list-style-type: none"> <li>– How much storage you have for instances</li> <li>– How much storage is being used for devices</li> <li>– What percentage of storage for a device is free for expansion</li> <li>– How much storage is being used for fragments</li> <li>– What percentage of storage for a fragment is free for expansion</li> </ul> </li> <li>• <b>UDB databases</b> <ul style="list-style-type: none"> <li>– How much allocated storage you have for a UDB Instance and the objects it contains</li> <li>– How much storage is being used for a UDB Instance and the objects it contains</li> <li>– What percentage of UDB storage is free for expansion</li> </ul> </li> </ul> <p>The Capacity reporting function provides you with the ability to view the storage capacity, used space, and free space for your RDBMS objects in a single view.</p> | Probes                                           |

| Report Category         | Use these reports to do the following tasks                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | What job collects information for these reports? |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------|
| Usage Reports           | <p>Monitor and report on the usage and growth of the consumption of your DBMS storage:</p> <ul style="list-style-type: none"> <li>• <b>Storage usage within all DBMSs.</b> If you are monitoring the multiple DBMSs, use this node to view common storage usage information for databases, table spaces, tables, and users across all monitored DBMSs.</li> <li>• <b>Oracle databases</b> <ul style="list-style-type: none"> <li>– View detailed information for table spaces including their total size, used space, empty used space, # of tables and indexes, and so on.</li> <li>– Identify the tables consuming the most space within your environment.</li> <li>– Identify the segments with the most extents and most empty used space.</li> <li>– Associate users with the amount of storage they are consuming on specific databases, table spaces, groups of table spaces, specific computers, groups of computers, and throughout the entire network.</li> </ul> </li> <li>• <b>SQL Server databases</b> <ul style="list-style-type: none"> <li>– View detailed information for databases including their total size, used space, empty used space, # of tables and indexes, and so on.</li> <li>– Identify the tables consuming the most space within your environment.</li> <li>– Associate users with the amount of storage they are consuming on specific databases, groups of databases, specific computers, groups of computers, and throughout the entire network.</li> </ul> </li> <li>• <b>Sybase databases</b> <ul style="list-style-type: none"> <li>– View detailed information for databases including their total size, used space, empty used space, # of tables and indexes, and so on.</li> <li>– Identify the tables consuming the most space within your environment.</li> <li>– Associate users with the amount of storage they are consuming on specific databases, groups of databases, specific computers, groups of computers, and throughout the entire network.</li> </ul> </li> <li>• <b>UDB databases</b> <ul style="list-style-type: none"> <li>– View detailed information for table spaces including their total size, used space, empty used space, # of tables and indexes, and so on.</li> <li>– Identify the tables consuming the most space within your environment. You can view this information at the table space, database, and computer level, as well as at a network-wide level.</li> <li>– Associate users with the amount of storage they are consuming on specific databases, table spaces, groups of table spaces, specific computers, groups of computers, and throughout the entire network.</li> </ul> </li> </ul> | Scans                                            |
| Usage Violation Reports | Create and enforce corporate storage usage policies. You can report on violations of quotas that are set at the network-wide, instance, table space, and database levels.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Scans, Quotas                                    |

### Reporting on RDBMS Objects:

Learn about the RDBMS objects upon which you can report using Data Manager for Databases reports.

The following table lists the RDBMS objects whose storage usage is monitored by Data Manager for Databases.

| Oracle         | SQL Server | Sybase    | DB2/UDB     |
|----------------|------------|-----------|-------------|
| Instances      | Instances  | Instances | Instances   |
| Databases      | Databases  | Devices   | Databases   |
| Tablespaces    | Datafiles  | Databases | Tablespaces |
| Datafiles      | Tables     | Fragments | Containers  |
| Tables         | Indexes    | Tables    | Tables      |
| Indexes        | Log Files  | Indexes   | Indexes     |
| Extents        |            |           |             |
| Segments       |            |           |             |
| Redo Log files |            |           |             |
| Control Files  |            |           |             |
| Archive Log    |            |           |             |
| Directories    |            |           |             |

**Note:**

1. Each of the major reporting categories in the Reporting facility can be expanded in the navigation tree to reveal additional sub-categories of reporting. Some of the nodes within the reporting tree will appear only if you set up access for the RDBMS to which that node applies. For example, if you set up for monitoring Oracle and SQL Server instances only, nodes representing the other RDBMSs supported by Data Manager for Databases (such as Sybase and UDB) will not appear within the navigation tree. If you have set up access for only one RDBMS, the **All DBMSs** node will also not be available in some nodes.
2. Data Manager for Databases does not currently support the monitoring of clustered database applications.

**Using By clauses:**

Use the By clauses associated with Data Manager for Databases reports as a method for viewing the data in those reports from different perspectives.

As you drill down through Data Manager for Databases's reports, you will notice a set of options that repeats for each report type. These *By Clauses* allow you to generate variations of the reports with different perspectives and groupings of the data. Some of the available By Clauses are described in the following table:

| By Clause                      | Description                                                                                                                                |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| <b>By DBMS Type</b>            | Use this clause to select the DBMS that you want to report upon.                                                                           |
| <b>By Instance</b>             | Use this clause to select the instances that you want to report upon and the report that will be generated based upon those instances.     |
| <b>By Device (Sybase only)</b> | Use this clause to select the devices that you want to report upon and the report that will be generated based upon those devices.         |
| <b>By Database</b>             | Use this clause to select the database that you want to report upon and the report that will be generated based upon those databases.      |
| <b>By Tablespace</b>           | Use this clause to select the tablespaces that you want to report upon and the report that will be generated based upon those tablespaces. |

| By Clause                         | Description                                                                                                                                                                       |
|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>By Tablespace Group</b>        | Use this clause to select the tablespace groups that you want to report upon and the report that will be generated based upon those tablespace groups.                            |
| <b>By Fragments (Sybase only)</b> | Use this clause to select the fragments that you want to report upon and the report that will be generated based upon those fragments.                                            |
| <b>By Container (UDB only)</b>    | Use this clause to select the containers that you want to report upon and the report that will be generated based upon those containers.                                          |
| <b>By Computer</b>                | Use this clause to select the computers that you want to report upon and the report that will be generated based upon those computers.                                            |
| <b>By Computer Group</b>          | Use this clause to select the computer groups that you want to report upon and the report that will be generated based upon those computer groups.                                |
| <b>By User</b>                    | Use this clause to select the users that you want to report upon and the report that will be generated based upon those users.                                                    |
| <b>Network Wide</b>               | The reports and charts that you see will be generated on all data, network wide. This allows you to get a global view of the storage and storage usage for your whole enterprise. |

## Asset reports

Use Asset reports to view information about the instances, devices, databases, tablespaces, tables, containers, data files, fragments, control files, redo logs, and archive log directories, and general hardware inventory of RDBMS storage assets within your environment.

While running probes and scans, Data Manager for Databases itemizes your RDBMS storage assets and provides a hierarchical view of that information so you can drill down to view your those assets in greater detail. You can view information about RDBMS storage resources through the following views:

| Asset Reporting View           | Use these reports to...                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Asset &gt; By Computer</b>  | View the instances and their related objects based upon the computers on which they reside.                                                                                                                                                                                                                                                                |
| <b>Asset &gt; By DBMS Type</b> | This node appears if you are using Data Manager for Databases to monitor more than one RDBMS. Expand this node to view asset information about instances organized according to their DBMS.                                                                                                                                                                |
| <b>Asset &gt; System-wide</b>  | Investigate all of your instances, databases, tablespaces, tables, container, data files, users and other assets for your <i>entire</i> system.<br><b>Note:</b> (UDB only) For the <b>All DBMSs</b> reports available under the <b>Asset &gt; System-wide</b> reporting nodes, the Database-Tablespaces view shows information about UDB tablespaces only. |

### By Computer, By DBMS Type:

Learn about the reports available within the **By Computer** and **By DBMS Type** nodes within Asset reporting.

| Asset Report Type                                                                          | Use these reports to...                                                                                                                                | What's in these reports?                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>By Computer</b><br>By Computer > <i>machine_name</i> ><br><i>instance/database_name</i> | Drill down through your DBMS storage assets in a hierarchical manner according to the computers on which they reside or the DBMS to which they belong. | Depending on the DBMS of the Instance you want to view, you can drill down and expose its: <ul style="list-style-type: none"> <li>• Tablespaces</li> <li>• Databases</li> <li>• Devices</li> <li>• Monitored tables</li> <li>• Containers</li> <li>• Data files</li> <li>• Fragments</li> <li>• Control files</li> <li>• Redo logs</li> <li>• Archive log directories</li> </ul> |
| <b>By DBMS Type</b><br>By DBMS Type > <i>dbms_type</i> ><br><i>instance/database_name</i>  |                                                                                                                                                        |                                                                                                                                                                                                                                                                                                                                                                                  |



| Asset Report Type                                                                                                                                                                                                    | Use these reports to...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | What's in these reports?                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Instance Information</b><br><br><b>By Computer</b><br><br>By Computer > <i>machine_name</i> ><br><i>instance_name</i><br><br><b>By DBMS Type</b><br><br>By DBMS Type > <i>dbms_type</i> ><br><i>instance_name</i> | View information about all of your monitored instances. Specifically, you can view information about the following information based upon DBMS type:<br><br><b>SQL Server</b> <ul style="list-style-type: none"> <li>• Databases</li> <li>• Monitored Tables</li> <li>• Data Files</li> </ul> <b>Oracle</b> <ul style="list-style-type: none"> <li>• Tablespaces</li> <li>• Monitored Tables</li> <li>• Data Files</li> <li>• Control Files</li> <li>• Redo Logs</li> <li>• Archive Log Directories</li> </ul> <b>Sybase</b> <ul style="list-style-type: none"> <li>• Devices</li> <li>• Databases</li> <li>• Monitored Tables</li> <li>• Fragments</li> </ul> <b>UDB</b> <ul style="list-style-type: none"> <li>• Instances</li> <li>• Databases</li> <li>• Tablespaces</li> <li>• Tables</li> <li>• Containers</li> <li>• Users</li> </ul> | <b>SQL Server</b> <ul style="list-style-type: none"> <li>• RDBMS Type, Name, Version, Loginid</li> <li>• Number of Databases, Number of Datafiles, Number of Log Files</li> <li>• Total Size of all Storage, Capacity of Data Files, Free Space of Data Files, Capacity of Log Files, Free Space of Log Files</li> <li>• Probe Status, Last Probe Time</li> </ul> <b>Oracle</b><br><br><i>Database Information tab</i> <ul style="list-style-type: none"> <li>• Name, Owner, Block Size, Log Mode, Status</li> <li>• Maximum Number of Instances, Number of Tablespaces, Maximum Number of Data Files, Number of Data Files, Maximum Number of Log Files, Number of Log Files</li> <li>• Total Size of all Storage, Capacity Of Data Files, Free Space of Data Files</li> <li>• Create Time, Discovered Time, Dropped Time</li> </ul> <i>Instance Information tab</i> <ul style="list-style-type: none"> <li>• RDBMS Type, Name, Version, Home Directory, Loginid, Port, Probe Status</li> <li>• Last Probe Time, Last Start Time, Discovered Time</li> </ul> <b>Sybase</b> <ul style="list-style-type: none"> <li>• RDBMS Type, Name, Version, Home Directory, Loginid, Port</li> <li>• Number of Databases, Number of Devices, Number of Fragments, Number of Log Fragments</li> <li>• Total Size of all Storage, Mirror Overhead, Capacity of Data Fragments, Free Space in Data Fragments, Capacity of Log Fragments, Free Space in Log Fragments</li> <li>• Probe Status, Last Probe Time</li> </ul> <b>UDB</b> <ul style="list-style-type: none"> <li>• RDBMS Type, Name, Version, Home Directory, Loginid, Port</li> <li>• Number of Tablespaces, Number of Containers, Number of Log Files</li> <li>• Total Size of all Storage, Container Capacity, DMS Container Free Space, Total Size of all Log Files</li> <li>• Probe Status, Last Probe Time</li> </ul> |

| Asset Report Type                                                                                                                                                                                                             | Use these reports to...                                                                   | What's in these reports?                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>Oracle Detailed Instance Information</b></p> <p><b>By Computer</b></p> <p>By Computer &gt; <i>machine_name</i> &gt; <i>instance_name</i></p>                                                                            | <p>View information about the objects associated with a selected Oracle Instance.</p>     | <ul style="list-style-type: none"> <li>• <b>Tablespaces:</b> When you expand the <b>Tablespaces</b> node, click the name of a tablespace to view information about that tablespace and its related data file.</li> <li>• <b>Monitored Tables:</b> When you expand the <b>Monitored Tables</b> node, click the name of a table to view its information. Only the tables that belong to Table Groups included in scans will be available to view through this node.</li> <li>• <b>Data Files:</b> Click the <b>Data Files</b> node, click the name of a data file to view information about the data files associated with the Instance, including: Data File Path, Used Space, Percent Used, Free Space, Total Size, Coalesced Extents, Minimum Free Extent Size, Maximum Free Extent Size, and Create Time.</li> <li>• <b>Control Files:</b> Click this node to view information on the control files associated with the Instance, including the File Name (full path) and File Size.</li> <li>• <b>Redo Logs:</b> Click this node to view information on the redo logs associated with the Instance. <ul style="list-style-type: none"> <li>– for Log Group: Sequence Number, Status, Archived, Total Size</li> <li>– for Log File: File Name (full path)</li> </ul> </li> <li>• <b>Archive Log Directories:</b> Click this node to view information on the archive log directory associated with the Instance, including Directory Name, Size of Files, File Count, Status, Last Failure Time.</li> </ul> |
| <p><b>SQL Server Detailed Instance Information</b></p> <p><b>By Computer</b> By Computer &gt; <i>machine_name</i> &gt; <i>server_name</i></p> <p><b>By DBMS Type</b> By DBMS Type &gt; SQL/Server &gt; <i>server_name</i></p> | <p>View information about the objects associated with a selected SQL Server Instance.</p> | <ul style="list-style-type: none"> <li>• <b>Databases:</b> When you expand the <b>Databases</b> node, click the name of a Database to view information about that database, its related data file, and its related log file.</li> <li>• <b>Monitored Tables:</b> When you expand the <b>Monitored Tables</b> node, click the name of a table to view its information including the computer, instances, and database on which is resides, as well as its creator, total size, empty used space, index size, number of indexes, row count, etc.</li> </ul> <p>Only the tables that belong to table groups included in scans will be available to view through this node.</p> <ul style="list-style-type: none"> <li>• <b>Data Files:</b> Click the <b>Data Files</b> node, click the name of a data file to view information about the data files associated with the Instance, including: Data File Path, Used Space, Percent Used, Free Space, Total Size, Coalesced Extents, Minimum Free Extent Size, Maximum Free Extent Size, and Create Time.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                               |

| Asset Report Type                                                                                                                                                                                    | Use these reports to...                                                        | What's in these reports?                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Sybase Detailed Instance Information</b><br><b>By Computer</b><br>By Computer > <i>machine_name</i> > <i>instance_name</i><br><b>By DBMS Type</b><br>By DBMS Type > Sybase > <i>instance_name</i> | view information about the objects associated with a selected Sybase Instance. | <ul style="list-style-type: none"> <li>• <b>Devices:</b> Click this node to view information about the Device associated with the Instance, including its name, computer on which it resides, path, free space, percent free, used space, total size, fragments, mirror path, etc.</li> <li>• <b>Databases:</b> When you expand the <b>Databases</b> node, click the name of a Database to view information about that database, its device information, and its related fragment information.</li> <li>• <b>Monitored Tables:</b> When you expand the <b>Monitored Tables</b> node, click the name of a table to view its information including the computer, instances, and database on which it resides, as well as its creator, total size, empty used space, index size, number of indexes, row count, etc.</li> </ul> <p>Only the tables that belong to table groups included in scans will be available to view through this node.</p> <ul style="list-style-type: none"> <li>• <b>Fragments:</b> Click this node to view about the Fragments associated with the Instance including the computer, instances, and database on which they resides as well as their path, used space, percent used, free space, total size, data type, etc.</li> </ul> |
| <b>UDB Detailed Instance Information</b><br><b>By Computer</b><br>By Computer > <i>machine_name</i> > <i>instance_name</i><br><b>By DBMS Type &gt;</b>                                               | view information about the objects associated with a selected UDB Instance.    | <ul style="list-style-type: none"> <li>• <b>Databases:</b> When you expand the <b>Databases</b> node, click the name of a databases to view information about that database.</li> <li>• <b>Monitored Tables:</b> When you expand the <b>Monitored Tables</b> node, click the name of a table to view its information. Only the tables that belong to Table Groups included in scans will be available to view through this node.</li> <li>• <b>Containers:</b> Click the <b>Containers</b> node, click the name of a container to view information about the containers associated with the Instance, including: Database, Tablespace, Container Node, Container Name, Used Space, Percent Used, Free Space, Total Size, Container Type, Status, Discovered Time.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

### System-wide:

Use System-wide reports to view your RDBMS storage assets as a whole, across the storage environment

| Asset Report Type                                   | Use these reports to...                                                                                       | What's in these reports?                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>System-Wide</b><br>Asset > System-wide           | drill down through your DBMS storage assets in a hierarchical manner.                                         | Depending on the DBMS of the Instance you want to view, you can drill down and expose its: <ul style="list-style-type: none"> <li>• Tablespaces</li> <li>• Databases</li> <li>• Devices</li> <li>• Monitored tables</li> <li>• Containers</li> <li>• Data files</li> <li>• Fragments</li> <li>• Control files</li> <li>• Redo logs</li> <li>• Archive log directories</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>All DBMSs</b><br>Asset > System-wide > All DBMSs | view information about all the DBMSs within your environment based on the objects related to those instances. | <ul style="list-style-type: none"> <li>• <b>Instances:</b> Expand the <b>Instances</b> node and click the sorting method by which you want to view instances.</li> <li>• <b>Databases-Tablespaces:</b> Expand <b>Databases-Tablespaces</b> node and click the sorting method by which you want to view information about the databases and/or tablespaces within your environment.</li> <li>• <b>Tables:</b> Expand the <b>Tables</b> node and click the sorting method by which you want to view tables. Only the tables that belong to table groups included in scans will appear within the reports viewed through this node.</li> <li>• <b>Files:</b> Expand the <b>Files</b> node and click the sorting method by which you want to view information about data files/containers, which includes: File Path, Free Space, Percent Free, Used Space, Total Size, Create Time, and Discovered Time.</li> <li>• <b>Users:</b> Expand the <b>Users</b> node and click the sorting method by which you want to view the users who are consuming space on the monitored DBMS within your environment.</li> </ul> |

---

**Asset Report Type****All DBMSs (continued)****What's in these reports?**

- **Instances:** Expand the **Instances** node and click the sorting method by which you want to view instances.
    - By Instance: sort based on Instance name
    - By Version: sort based on the version of the DBMS under which an Instance is running
    - By Start Time: sort based upon when an Instance was last started
    - By Probe Time: sort based upon when an Instance was last probed
    - By RDBMS Type: sort based upon the DBMS of an Instance
  - **Databases-Tablespaces:** Expand the **Databases-Tablespaces** node and click the sorting method by which you want to view information about the databases or tablespaces within your environment.
    - By Database-Tablespace: sort based on DB/TB name
    - By Freespace: sort based on the amount of free space on a DB/TB
    - By Total Size: sort based on the total size of a DB/TB
    - By Table Count: sort based upon the number of tables on a DB/TB
    - By Scan Time: sort based upon when a DB/TB was last scanned
    - By Discovered Time: sort based upon when a DB/TB was discovered
    - Dropped Tablespaces: sort based upon dropped DBs or TBs.
    - By RDBMS Type: sort based upon the DBMS of a DB/TB
  - **Tables:** Expand the **Tables** node and click the sorting method by which you want to view tables. Only the tables that belong to Table Groups included in scans will appear within the reports viewed through this node.
    - By Table: sort based on table name
    - By Number of Indexes: sort based upon the number of indexes associated with a table
    - By Total Size: sort based on the total size of a table
    - By Row Count: sort based upon the number of rows in a table
    - By Chained Row Count: sort based upon the number of chained rows in a table
    - By Discovered Time: sort based upon when a table was discovered
    - Dropped Tables: sort based upon dropped tables
    - By RDBMS Type: sort based upon the DBMS of a table
  - **Files:** Expand the **Files** node and click the sorting method by which you want to view information about DBMS files.
    - By File: sort based on file name
    - By Free Space: sort based upon file free space
    - By Total Size: sort based on the total size of a file
    - By Discovered Time: sort based upon when a file was discovered
    - By RDBMS Type: sort based upon the DBMS of a file
  - **Users:** Expand the **Users** node and click the sorting method by which you want to view the users who are consuming space on the monitored DBMS within your environment.
-

| Asset Report Type                             | What's in these reports?                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Oracle</b><br>Asset > System-wide > Oracle | <ul style="list-style-type: none"> <li>• <b>Instances:</b> Expand the <b>Instances</b> node and click the sorting method by which you want to view instances.               <ul style="list-style-type: none"> <li>– By Instance: sort based on Instance name</li> <li>– By Version: sort based on the version of the DBMS under which an Instance is running</li> <li>– By Start Time: sort based upon when an Instance was last started</li> <li>– By Probe Time: sort based upon when an Instance was last probed</li> </ul> </li> <li>• <b>Tablespaces:</b> Expand this node and click the sorting method by which you want to view information about tablespaces within your environment.               <ul style="list-style-type: none"> <li>– By Tablespace: sort based on tablespace name</li> <li>– By Max Free Extent: sort based on the maximum free extent size</li> <li>– By Freespace: sort based on the amount of free space on a tablespace</li> <li>– By Total Size: sort based on the total size of a tablespace</li> <li>– By Coalesced Extents: sort based on the number of coalesced extents</li> <li>– By Table Count: sort based upon the number of tables</li> <li>– By Scan Time: sort based upon when a tablespace was last scanned</li> <li>– By Discovered Time: sort based upon when a tablespace was discovered</li> <li>– Dropped Tablespaces: sort based upon dropped tablespaces</li> </ul> </li> <li>• <b>Tables:</b> Expand the <b>Tables</b> node and click the sorting method by which you want to view tables. Only the tables that belong to Table Groups included in scans will appear within the reports viewed through this node.               <ul style="list-style-type: none"> <li>– By Table: sort based on table name</li> <li>– By Number of Indexes: sort based upon the number of indexes associated with a table</li> <li>– By Total Size: sort based on the total size of a table</li> <li>– By Row Count: sort based upon the number of rows in a table</li> <li>– By Chained Row Count: sort based upon the number of chained rows in a table</li> <li>– By Discovered Time: sort based upon when a table was discovered</li> <li>– Dropped Tables: sort based upon dropped tables</li> </ul> </li> <li>• <b>Data Files:</b> Expand the <b>Data Files</b> node and click the sorting method by which you want to view information about DBMS files.               <ul style="list-style-type: none"> <li>– By Data File: sort based on file name</li> <li>– By Max Free Extent: sort data files sorted by their max free extent sizes</li> <li>– By Free Space: sort based upon file free space</li> <li>– By Total Size: sort based on the total size of a file</li> <li>– With Auto-extend: sort data files that have auto-extend specified</li> <li>– Without Auto-extend: sort data files that do not have auto-extend specified</li> <li>– By Discovered Time: sort based upon when a file was discovered</li> </ul> </li> <li>• <b>Users:</b> Expand the <b>Users</b> node and click the sorting method by which you want to view the users who are consuming space on the monitored Oracle instances within your environment.</li> </ul> |

---

**Asset Report Type****SQL Server**

Asset > System-wide >  
SQL/Server

---

**What's in these reports?**

- **Instances:** Expand the **Instances** node and click the sorting method by which you want to view instances.
    - By Instance: sort based on Instance name
    - By Version: sort based on the version of the DBMS under which an Instance is running
    - By Probe Time: sort based upon when an Instance was last probed
  - **Databases:** Expand this node and click the sorting method by which you want to view information about databases within your environment.
    - By Database: sort based on database name
    - By Freespace: sort based on the amount of free space on a database
    - By Total Size: sort based on the total size of a database
    - By Table Count: sort based upon the number of tables in the database
    - By Scan Time: sort based upon when a database was last scanned
    - By Discovered Time: sort based upon when a database was discovered
    - Dropped Databases: sort based upon dropped databases
  - **Tables:** Expand the **Tables** node and click the sorting method by which you want to view tables. Only the tables that belong to Table Groups included in scans will appear within the reports viewed through this node.
    - By Table: sort based on table name
    - By Number of Indexes: sort based upon the number of indexes associated with a table
    - By Total Size: sort based on the total size of a table
    - By Row Count: sort based upon the number of rows in a table
    - By Discovered Time: sort based upon when a table was discovered
    - Dropped Tables: sort based upon dropped tables
  - **Data Files:** Expand the **Data Files** node and click the sorting method by which you want to view information about data files.
    - By Data File: sort based on file name
    - By Free Space: sort based upon file free space
    - By Total Size: sort based on the total size of a file
    - With Auto-extend: sort data files that have auto-extend specified
    - Without Auto-extend: sort data files that do not have auto-extend specified
    - By Discovered Time: sort based upon when a file was discovered
  - **Users:** Expand the **Users** node and click the sorting method by which you want to view the users who are consuming space on the monitored SQL Server instances.
    - By Users: sort based on user names
    - By Discovered Time: sort based when a user was discovered
-

| Asset Report Type | What's in these reports?                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Sybase            | <ul style="list-style-type: none"> <li>• <b>Instances:</b> Expand the <b>Instances</b> node and click the sorting method by which you want to view instances. <ul style="list-style-type: none"> <li>– By Instance: sort based on Instance name</li> <li>– By Version: sort based on the version of the DBMS under which an Instance is running</li> <li>– By Probe Time: sort based upon when an Instance was last probed</li> </ul> </li> <li>• <b>Devices:</b> Expand this node and click the sorting method by which you want to view information about devices within your environment. <ul style="list-style-type: none"> <li>– By Device: sort based on device name</li> <li>– By Freespace: sort based on the amount of free space on a device</li> <li>– By Total Size: sort based on the total size of a device</li> <li>– By Fragments: sort based upon the number of fragments for a device</li> <li>– By Discovered Time: sort based upon when a device was discovered</li> <li>– Dropped Devices: sort based upon dropped devices</li> </ul> </li> <li>• <b>Databases:</b> Expand this node and click the sorting method by which you want to view information about databases within your environment. <ul style="list-style-type: none"> <li>– By Database: sort based on database name</li> <li>– By Freespace: sort based on the amount of free space on a database</li> <li>– By Total Size: sort based on the total size of a database</li> <li>– By Table Count: sort based upon the number of tables in the database</li> <li>– By Scan Time: sort based upon when a database was last scanned</li> <li>– By Discovered Time: sort based upon when a database was discovered</li> <li>– By Last Dump Time: sort based on the date/time when devices were last dumped</li> <li>– Dropped Databases: sort based upon dropped databases</li> </ul> </li> <li>• <b>Fragments:</b> Expand this node and click the sorting method by which you want to view information about fragments. <ul style="list-style-type: none"> <li>– By Fragments: sort based on fragment name</li> <li>– By Free Space: sort based upon the free space in fragments</li> <li>– By Total Size: sort based on the total size of fragments</li> <li>– By Discovered Time: sort based upon when fragments were discovered</li> </ul> </li> <li>• <b>Tables:</b> Expand the <b>Tables</b> node and click the sorting method by which you want to view tables. Only the tables that belong to Table Groups included in scans will appear within the reports viewed through this node. <ul style="list-style-type: none"> <li>– By Table: sort based on table name</li> <li>– By Number of Indexes: sort based upon the number of indexes associated with a table</li> <li>– By Total Size: sort based on the total size of a table</li> <li>– By Row Count: sort based upon the number of rows in a table</li> <li>– By Chained Row Count: sort based upon the number of chained rows in a table</li> <li>– By Discovered Time: sort based upon when a table was discovered</li> <li>– Dropped Tables: sort based upon dropped tables</li> </ul> </li> <li>• <b>Users:</b> Expand the <b>Users</b> node and click the sorting method by which you want to view the users who are consuming space on the monitored Sybase instances.</li> </ul> |



---

**Asset Report Type****UDB**

Asset &gt; System-wide &gt; UDB

**What's in these reports?**

- **Instances:** Expand the **Instances** node and click the sorting method by which you want to view instances.
    - By Instance: sort based on Instance name
    - By Version: sort based on the version of the DBMS under which an Instance is running
    - By Probe Time: sort based upon when an Instance was last probed
  - **Databases:** Expand this node and click the sorting method by which you want to view information about databases within your environment.
    - By Database: sort based on database name
    - By Freespace: sort based on the amount of free space on a database
    - By Total Size: sort based on the total size of a database
    - By Table Count: sort based upon the number of tables in the database
    - By Scan Time: sort based upon when a database was last scanned
    - By Discovered Time: sort based upon when a database was discovered
    - Dropped Databases: sort based upon dropped databases
  - **Tablespaces:** Expand this node and click the sorting method by which you want to view information about tablespaces.
    - By Tablespace: sort based on tablespace name
    - By FreeSpace: sort based upon the freespace in tablespaces
    - By Total Size: sort based on the total size of tablespaces
    - By Table Count: sort based upon the number of tables associated with tablespaces
    - By Scan Time: sort based upon the date/time when tablespaces were scanned
    - By Discovered Time: sort based upon the date/time when tablespaces were discovered
    - Dropped Tablespaces: sort based upon the names of dropped tablespaces
  - **Tables:** Expand the **Tables** node and click the sorting method by which you want to view tables. Only the tables that belong to Table Groups included in scans will appear within the reports viewed through this node.
    - By Table: sort based on table name
    - By Number of Indexes: sort based upon the number of indexes associated with a table
    - By Total Size: sort based on the total size of a table
    - By Row Count: sort based upon the number of rows in a table
    - By Chained Row Count: sort based upon the number of chained rows in a table
    - By Discovered Time: sort based upon the date/time when a table was discovered
    - Dropped Tables: sort based upon the names of dropped tables
-

| Asset Report Type | What's in these reports?                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                   | <ul style="list-style-type: none"> <li>• <b>Containers:</b> Expand the <b>Containers</b> node and click the sorting method by which you want to view containers. <ul style="list-style-type: none"> <li>– By Container: sort based on container name</li> <li>– By Freespace: sort based upon the freespace in containers</li> <li>– By Total Size: sort based on the total size of containers</li> <li>– With Auto-extend: sort based upon whether containers with Auto-extend on</li> <li>– Without Auto-extend: sort based upon whether containers with Auto-extend off</li> <li>– By Discovered Time: sort based upon the date/time when container were discovered</li> </ul> </li> <li>• <b>Users:</b> Expand the <b>Users</b> node and click the sorting method by which you want to view the users who are consuming space on the monitored UDB instances.</li> </ul> |

## Capacity reports

Use Capacity reports to view the storage capacity, used space, and free space for your RDBMS objects.

Capacity enable you to view storage capacity at the computer, instance, database, device, fragment, container, and data file level. For a specific Instance, computer, group of computers, database, group of databases, tablespace, group of tablespaces, datafiles, containers, fragments, or all the computers in your enterprise, you can view the following:

- How much storage you have for a DBMS and the objects it contains
- How much storage is being used for a DBMS and the objects it contains
- What percentage of DBMS storage is free for expansion

### All DBMSs Capacity reports:

Use these reports to view storage capacity information about the different RDBMS instances that you are monitoring within your environment.

| Capacity Reporting View                       | Use these reports to...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Capacity > All DBMSs > Total Instance Storage | <p>If you have multiple RDBMSs that you monitor using Data Manager, use this node to view detailed storage information about the different RDBMS instances within your environment.</p> <p>Depending on the RDBMS of the instance you want to view, use this report to view information about its:</p> <ul style="list-style-type: none"> <li>• RDBMS instance</li> <li>• RDBMS type</li> <li>• total size</li> <li>• file capacity</li> <li>• file free space</li> <li>• DB-TS count</li> <li>• file count</li> <li>• log file count</li> </ul> |

### Oracle Capacity reports:

Use these reports to view storage capacity information about the Oracle instances that you are monitoring within your environment.

| Capacity Reporting View                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Use these reports to...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>Oracle</b></p> <p>Capacity &gt; Oracle</p> <p><b>By Clauses</b> for these reports:</p> <ul style="list-style-type: none"> <li>• By Instance: sort based on Instance name</li> <li>• By Database: sort a report by database name</li> <li>• By Tablespace: sort a report by tablespace name</li> <li>• By Tablespace Group: view the desired storage information organized by tablespace group</li> <li>• By Datafile: sort based on file name</li> <li>• By Computer: sort based on the computer where the object's Instance resides</li> <li>• By Computer Group: view the desired storage information organized by computer group</li> <li>• Network-wide: view the desired storage information for your entire network</li> </ul> | <p>This node appears if you are using Data Manager for Databases to monitor Oracle instances. Expand this node to view the following information:</p> <ul style="list-style-type: none"> <li>• <b>Total Instance Storage.</b> Generate a report under this node to view detailed information about Oracle instances, including: <ul style="list-style-type: none"> <li>– total size</li> <li>– file capacity</li> <li>– file free space</li> <li>– DB-TS count</li> <li>– file count</li> <li>– log file count</li> </ul> </li> <li>• <b>Data File Used Space.</b> Generate a report under this node to view detailed information about the storage usage of datafiles associated with the tablespaces in an Instance, including: <ul style="list-style-type: none"> <li>– tablespace</li> <li>– data file path</li> <li>– used space</li> <li>– percent used</li> <li>– free space</li> <li>– total size</li> <li>– free extents</li> <li>– coalesced extents</li> <li>– maximum free extent size</li> </ul> </li> <li>• <b>Data File Free Space.</b> Generate a report under this node to view detailed information about the storage free space of datafiles associated with the tablespaces in an Instance, including: <ul style="list-style-type: none"> <li>– tablespace</li> <li>– data free space</li> <li>– percent free</li> <li>– data used space</li> <li>– data size</li> <li>– free extents</li> <li>– coalesced extents</li> <li>– maximum free extent size</li> </ul> </li> </ul> |

### SQL Server Capacity reports:

Use these reports to view storage capacity information about the SQL Server instances that you are monitoring within your environment.

---

## Capacity Reporting View

### SQL Server

Capacity > SQL/Server

#### By Clauses for these reports:

- By Instance: sort based on Instance name
- By Database: sort a report by database name
- By Database Group: view the desired storage information organized by database group
- By Datafile: sort based on file name
- By Computer: sort based on the computer where the object's Instance resides
- By Computer Group: view the desired storage information organized by computer group
- Network-wide: view the desired storage information for your entire network

## Use these reports to...

This node appears if you are using Data Manager for Databases to monitor SQL Server instances. Expand this node to view the following information:

- **Total Instance Storage.** Generate a report under this node to view detailed information about SQL Server instances, including:
    - total size
    - data file capacity
    - data file free space
    - log file capacity
    - log file free space
    - database count
    - data file count
    - log file count
  - **Data File Used Space.** Generate a report under this node to view detailed information about the storage usage of datafiles associated with the databases in an Instance, including:
    - database
    - data used space
    - percent used
    - data free space
    - data size
    - log used space
    - log percent used
    - log free space
    - log size
    - number data files
  - **Data File Free Space.** Generate a report under this node to view detailed information about the storage free space of datafiles associated with the databases in an Instance, including:
    - database
    - data free space
    - percent free
    - data used space
    - data size
    - log free space
    - log percent used
    - log used space
    - log size
    - number data files
- 

### Sybase Capacity reports:

Use these reports to view storage capacity information about the Sybase instances that you are monitoring within your environment.

---

## Capacity Reporting View

### Sybase

Capacity > Sybase

#### By Clauses for these reports:

- By Device: sort based on Device name
- By Instance: sort based on Instance name
- By Fragments: sort based upon fragments
- By Database: sort a report by database name
- By Database Group: view the desired storage information organized by database group
- By Datafile: sort based on file name
- By Computer: sort based on the computer where the object's Instance resides
- By Computer Group: view the desired storage information organized by computer group
- Network-wide: view the desired storage information for your entire network

#### Use these reports to...

This node appears if you are using Data Manager for Databases to monitor Sybase instances. Expand this node to view the following information:

- **Total Instance Storage.** Generate a report under this node to view detailed information about Sybase instances, including:
    - total device size
    - mirrored device size
    - device free space
    - data fragment capacity
    - data fragment free space
    - log fragment capacity
    - log fragment free space
    - device count
    - database count
    - fragment count
    - log fragment count
  - **Device Used Space.** Generate a report under this node to view detailed information about the storage usage of devices associated with instances, including:
    - Instance
    - device path
    - used space
    - percent used
    - free space
    - total size
    - number fragments
    - mirror path
    - number devices
  - **Device Free Space.** Generate a report under this node to view detailed information about the storage free space of devices associated instances, including:
    - Instance
    - device path
    - free space
    - percent free
    - used space
    - total size
    - number fragments
    - low, high
    - mirror path
    - number devices
-

---

**Capacity Reporting View****Use these reports to...**

Sybase (continued)

- **Fragment Used Space.** Generate a report under this node to view detailed information about the storage usage of fragments associated with databases, including:
    - Instance
    - fragment path
    - database
    - data used space
    - percent used
    - data free space
    - data size
    - data type
    - log used space
    - log percent used
    - log free space
    - log size
    - number fragments
  - **Fragment Free Space.** Generate a report under this node to view detailed information about the storage free space of fragments associated databases, including:
    - Instance
    - fragment path
    - database
    - data free space
    - percent free
    - data used space
    - data size
    - log free space
    - log used space
    - log percent used
    - log used space
    - log size
    - number fragments
- 

**UDB Capacity reports:**

Use these reports to view storage capacity information about the UDB instances that you are monitoring within your environment.

---

## Capacity Reporting View

### UDB

Capacity > UDB

**By Clauses** for these reports:

- By Container: sort based on Container name
- By Tablespace: sort based on Tablespace name
- By Database: sort a report by database name
- By Database Group: view the desired storage information organized by database group
- By Instance: sort based on Instance name
- By Computer: sort based on the computer where the object's Instance resides
- By Computer Group: view the desired storage information organized by computer group
- Network-wide: view the desired storage information for your entire network

## Use these reports to...

This node appears if you are using Data Manager for Databases to monitor UDB instances. Expand this node to view the following information:

- **Total Instance Storage.** Generate a report under this node to view detailed information about UDB instances, including:
    - total size
    - container capacity
    - container free space
    - log file capacity
    - tablespace count
    - container count
    - log file count
    - log mode
  - **Container Used Space.** Generate a report under this node to view detailed information about the storage usage of containers associated with a tablespace, including:
    - Instance
    - database
    - tablespace
    - container node
    - container name
    - used space
    - percent used
    - free space
    - total size
    - container type
    - status
    - discovered time
  - **Container Free Space.** Generate a report under this node to view detailed information about the storage free space of containers associated with a tablespace, including:
    - Instance
    - database
    - tablespace
    - container node
    - container name
    - free space
    - percent free
    - used space
    - total size
    - container type
    - status
    - discovered time
- 

## Usage reports

Use Usage reports to monitor and report on the usage and growth of the storage consumption for your databases, tablespaces, tables, segments, and users.

Use these reports to do the following:

- View detailed information for tablespaces and databases including their total size, used space, empty used space, # of tables and indexes, etc.
- Identify the tables consuming the most space within your environment. You can view this information at the tablespace, database, and computer level, as well as at a network-wide level.
- (Oracle) Identify the segments with the most extents and most empty used space. You can view this information at the tablespace, database, and computer level, as well as at a network-wide level.
- Associate users with the amount of storage they are consuming on specific databases, groups of databases, tablespaces, groups of tablespaces, specific computers, groups of computers, and throughout the entire network.

#### All DBMS Usage reports:

Use these reports to view storage usage information about the different RDBMS instances that you are monitoring within your environment.

| Usage Reporting View                      | Use these reports to...                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Usage > All DBMSs                         | If you have multiple RDBMSs that you monitor using Data Manager, use this node to view detailed storage information about the different RDBMS instances within your environment.                                                                                                                                                                                                                                                                        |
| Usage > All DBMSs > Databases-Tablespaces | view storage usage statistics about the databases and tablespaces with your monitored DBMSs. Some of the information you can view through these reports includes: <ul style="list-style-type: none"><li>• database name, tablespace name</li><li>• computer</li><li>• Instance</li><li>• Instance type</li><li>• total size</li><li>• used space</li><li>• empty used space</li><li>• table count</li><li>• index count</li><li>• other count</li></ul> |
| Usage > All DBMSs > Tables                | view storage usage statistics about the tables with your monitored DBMSs. Some of the information you can view through these reports includes the following: <ul style="list-style-type: none"><li>• Largest tables</li><li>• Total size</li><li>• Creator</li><li>• Table name</li><li>• Empty used space</li><li>• Index size</li><li>• Row count</li><li>• Chained row count</li></ul>                                                               |



---

**Usage Reporting View**

Usage &gt; All DBMSs &gt; Users

**Use these reports to...**

associate users with the amount of storage they are consuming on specific DBMS objects, specific computers, groups of computers, and throughout the entire network.

- **Group by User Name:** select the users and then the resources on which you want to report storage consumption.

For example, when setting up a report to view user storage consumption on an Oracle or UDB tablespace, you must do the following:

- Select the users upon which you want to report
- Select the tablespaces whose storage consumption by the selected users you want to view

- **Group by Resource:** select the resources and then the users on which you want to report storage consumption.

For example, when setting up a report to view storage consumption on Oracle or UDB tablespaces for users, you must do the following:

- Select the tablespaces upon which you want to report
  - Select the users whose storage consumption on the selected tablespaces you want to view
- 

**Oracle Usage reports:**

Use these reports to view storage usage information about the Oracle tablespaces, tables, segments, and users that you are monitoring within your environment.

---

**Usage Reporting View****Oracle**

Usage &gt; Oracle

**By Clauses** for these reports:

- By Database—sort a report by database name
- By Tablespace—sort a report by tablespace name
- By Tablespace Group—view the desired storage information organized by tablespace group
- By Table—sort a report by table name
- By Table Group—view the desired storage information organized by table group
- By Computer—sort based on the computer where the object's Instance resides
- By Computer Group—view the desired storage information organized by computer group
- Network-wide—view the desired storage information for your entire network

**Use these reports to...**

This node appears if you are using Data Manager for Databases to monitor Oracle instances. Expand this node to view the following information:

- **Tablespaces.** View detailed information for tablespaces including the following:
    - Instance
    - Database
    - Total size
    - Used space
    - Empty used space
    - Freelist blocks
    - Table count
    - Index count
  - **Tables.** Identify the tables consuming the most space within your environment. You can view this information at the tablespace, database, and computer level, as well as at a network-wide level. Some of the information in these reports includes the following:
    - Table name
    - Total size
    - Empty used space
    - Index size
    - Freelist blocks
    - Number partitions
    - Number indexes
    - Number segments
    - Number extents
    - Percent free, used
-

---

**Usage Reporting View****Oracle**

Usage → Oracle

**By Clauses** for these reports:

- By Database—sort a report by database name
- By Tablespace—sort a report by tablespace name
- By Tablespace Group—view the desired storage information organized by tablespace group
- By Table—sort a report by table name
- By Table Group—view the desired storage information organized by table group
- By Computer—sort based on the computer where the object's Instance resides
- By Computer Group—view the desired storage information organized by computer group
- Network-wide—view the desired storage information for your entire network

**Use these reports to...**

This node appears if you are using Data Manager for Databases to monitor Oracle instances. Expand this node to view the following information:

- **Segments.** Identify the segments with the most extents and most empty used space. Some of the information in these reports includes the following:
    - Number extents
    - Segment name, creator
    - Tablespace
    - Parent name, creator
    - Total size
    - Empty used space
    - Number extents
    - Freelist blocks, groups
  - **Users.** Associate users with the amount of storage they are consuming on specific databases, tablespaces, groups of tablespaces, specific computers, groups of computers, throughout the entire network, and applied Quotas. Some of the information included in these reports includes the following:
    - Computer, instance, database, tablespace
    - Total size
    - Empty used space
    - Table count
    - Index count
    - Largest segment, 2nd largest segment
    - Segment sizes: <50Kb, 500Kb-100Mb, 1Mb-10Mb, 10Mb-100Mb, 100Mb-500MB, 500Mb-1Gb, >1Gb
    - Quota name
    - Percent of Quota
    - Quota threshold
- 

**SQL Server Usage reports:**

Use these reports to view storage usage information about the SQL Server databases, tables, and users that you are monitoring within your environment.

---

## Usage Reporting View

### SQL Server

Usage > SQL/Server

**By Clauses** for these reports:

- By Instance: sort based on Instance name
- By Database: sort a report by database name
- By Database Group: view the desired storage information organized by database group
- By Computer: sort based on the computer where the object's Instance resides
- By Computer Group: view the desired storage information organized by computer group
- Network-wide: view the desired storage information for your entire network
- By User Quota: sort based upon the users to which Quotas are applied. A user who has not violated a quota will not appear in this report.

## Use these reports to...

This node appears if you are using Data Manager for Databases to monitor storage usage on SQL Server instances. Expand this node to view the following information:

- **Databases.** Generate a report under this node to view detailed information about the storage on SQL Server databases, including the following:
    - Total size
    - Used space
    - Empty used space
    - Table count
    - Index count
  - **Tables.** Identify the tables consuming the most space within your environment. You can view this information at the table, database, and computer level, as well as at a network-wide level. Some of the information in these reports includes the following:
    - Table name
    - Total size
    - Empty used space
    - Index count
    - Row count
  - **Users.** Associate users with the amount of storage they are consuming on specific databases, groups of databases, specific computers, groups of computers, throughout the entire network, and applied Quotas . Some of the information in these reports includes the following:
    - Computer, instance, database
    - Total size
    - Empty used space
    - Table count
    - Index count
    - Largest object, 2nd largest object
    - Object sizes: <50Kb, 500Kb-100Mb, 1Mb-10Mb, 10Mb-100Mb, 100Mb-500MB, 500Mb-1Gb, >1Gb
    - Quota name
    - Percent of Quota
    - Quota threshold
- 

## Sybase Usage reports:

Use these reports to view storage usage information about the Sybase databases, tables, and users that you are monitoring within your environment.

---

## Usage Reporting View

### Sybase

Usage > Sybase

- By Database: sort a report by database name
- By Database Group: view the desired storage information organized by database group
- By Computer: sort based on the computer where the object's Instance resides
- By Computer Group: view the desired storage information organized by computer group
- Network-wide: view the desired storage information for your entire network
- By User Quota: sort based upon the users to which Quotas are applied. A user who has not violated a quota will not appear in this report.

## Use these reports to...

This node appears if you are using Data Manager for Databases to monitor storage usage on Sybase instances. Expand this node to view the following information:

- **Databases.** Generate a report under this node to view detailed information about the storage on Sybase databases, including the following:
    - Database name, computer, instance
    - Total size
    - Used space
    - Empty used space
    - Table count
    - Index count
  - **Tables.** Identify the tables consuming the most space within your environment. You can view this information at the table, database, and computer level, as well as at a network-wide level. Some of the information in these reports includes the following:
    - Table name, database, computer, Instance, creator
    - Total size
    - Empty used space
    - Index size
    - Maximum rows per page
    - Expected row size
    - Fill factor
    - Page gap
    - Average row length
    - Row count
  - **Users.** Associate users with the amount of storage they are consuming on specific databases, groups of databases, specific computers, groups of computers, throughout the entire network, and applied Quotas. Some of the information included in these reports includes the following:
    - Computer, instance, database
    - Total size
    - Empty used space
    - Table, index count
    - Largest object, 2nd largest object
    - Object sizes: <50Kb, 500Kb-100Mb, 1Mb-10Mb, 10Mb-100Mb, 100Mb-500MB, 500Mb-1Gb, >1Gb
    - Quota name
    - Percent of Quota
    - Quota threshold
- 

### UDB reports:

Use these reports to view storage usage information about the UDB tablespaces, tables, and users that you are monitoring within your environment.

---

## Usage Reporting View

### UDB

Usage > UDB

- By Tablespace: sort a report by tablespace name
- By Tablespace Group: view the desired storage information organized by tablespace group
- By Database: sort a report by database name
- By Instance: sort a report by Instance name
- By Computer: sort based on the computer where the object's Instance resides
- By Computer Group: view the desired storage information organized by computer group
- Network-wide: view the desired storage information for your entire network

## Use these reports to...

This node appears if you are using Data Manager for Databases to monitor storage usage on UDB instances. Expand this node to view the following information:

- **Tablespaces.** Generate a report under this node to view detailed information about the storage on UDB tablespaces, including the following:
  - Database name, computer, instance
  - Total size
  - Used space
  - Empty used space
  - Table count
  - Index count
- **Tables.** Identify the tables consuming the most space within your environment. You can view this information at the table, database, and computer level, as well as at a network-wide level. Some of the information in these reports includes the following:
  - Table name, database, computer, Instance, creator
  - Total size
  - Empty used space
  - Index size
  - Maximum rows per page
  - Expected row size
  - Fill factor
  - Page gap
  - Average row length
  - Row count

**Note:** When viewing reports under the **Monitored Tables** node, keep in mind that a table will appear in this report if: 1) it is included in the definition of a table group and 2) the table group in which it is included is defined as part of a scan.

- **Users.** Associate users with the amount of storage they are consuming on specific UDB instances, tablespaces, groups of tablespaces, computers, groups of computers, and throughout the entire network. Some of the information included in these reports includes the following:
  - Computer, instance
  - Total size
  - Empty used space
  - Table, index count
  - Largest object, 2nd largest object
  - Object sizes: <50Kb, 500Kb-100Mb, 1Mb-10Mb, 10Mb-100Mb, 100Mb-500MB, 500Mb-1Gb, >1Gb
  - Quota name
  - Percent of Quota
  - Quota threshold

---

## Usage Violation reports

Use Usage Violation reports to view information about quota violations and help enforce corporate storage usage policies.

With quotas, you can control how much storage a user, or a group of users, can consume on the following:

- Database or group of databases
- Tablespace or group of tablespaces
- Instance

- Throughout the entire network

The following table contains information about Usage Violation reports:

| Usage Violation Reporting Type                                                          | Use these reports to...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | What's in these reports?                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Quota Violations</b><br>Usage Violations > Quota Violations                          | identify the quota violations within the environment                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>All Quotas</b><br>Usage Violations > Quota Violations > All Quotas                   | view all the quotas and quota violations in the environment <ul style="list-style-type: none"> <li>• <b>By User.</b> View information about the users who have violated quotas.</li> <li>• <b>By User Quota.</b> View information about the quotas that have been violated.</li> </ul>                                                                                                                                                                                                                        | <ul style="list-style-type: none"> <li>• User name</li> <li>• RDBMS type</li> <li>• Instance</li> <li>• Database-tablespace</li> <li>• Space used</li> <li>• Threshold</li> <li>• % of threshold</li> <li>• Number of violations</li> <li>• Quota creator</li> <li>• Quota name</li> <li>• Quota type</li> </ul>                                                                                                                    |
| <b>Network Wide Quotas</b><br>Usage Violations > Quota Violations > Network-wide Quotas | view quota violations at the network-wide level <ul style="list-style-type: none"> <li>• <b>By User.</b> View information about the users who have violated quotas.</li> <li>• <b>By User Quota.</b> View information about the quotas that have been violated and which users have violated them.</li> </ul>                                                                                                                                                                                                 | <ul style="list-style-type: none"> <li>• User name</li> <li>• Space used</li> <li>• Threshold</li> <li>• % of threshold</li> <li>• Number of violations</li> <li>• Quota creator</li> <li>• Quota name</li> <li>• Quota type</li> </ul>                                                                                                                                                                                             |
| <b>Instance Quotas</b><br>Usage Violations > Quota Violations > Instance Quotas         | view quota violations at the Instance level <ul style="list-style-type: none"> <li>• <b>By Instance.</b> View Instance quota violations sorted by Instance</li> <li>• <b>By Computer.</b> View Instance quota violations sorted by computer</li> <li>• <b>By Computer Group.</b> View Instance quota violations sorted by computer group</li> <li>• <b>By User.</b> View Instance quota violations sorted by users</li> <li>• <b>By User Quota.</b> View Instance quota violations sorted by quota</li> </ul> | <ul style="list-style-type: none"> <li>• Computer</li> <li>• RDBMS type</li> <li>• Group creator (by group clauses only)</li> <li>• Group name (by group clauses only)</li> <li>• Instance</li> <li>• Database-tablespace</li> <li>• User name</li> <li>• Space used</li> <li>• Threshold</li> <li>• % of threshold</li> <li>• Number of violations</li> <li>• Quota creator</li> <li>• Quota name</li> <li>• Quota type</li> </ul> |

| Usage Violation Reporting Type                                                                        | Use these reports to...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | What's in these reports?                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Database-Tablespace Quotas</b><br>Usage Violations > Quota Violations > Database-Tablespace Quotas | view quota violations at the tablespace level <ul style="list-style-type: none"> <li>• <b>By Database-Tablespace.</b> View quota violations sorted by the database or tablespace name on which the violation occurred</li> <li>• <b>By Database-Tablespace Group.</b> View quota violations sorted by the database or tablespace group in which the violation occurred</li> <li>• <b>By Instance.</b> View database or tablespace quota violations sorted by Instance</li> <li>• <b>By Computer.</b> View database or tablespace quota violations sorted by computer</li> <li>• <b>By Computer Group.</b> View database or tablespace quota violations sorted by computer group</li> <li>• <b>By User.</b> View database or tablespace quota violations sorted by users</li> <li>• <b>By User Quota.</b> View database or tablespace quota violations sorted by quota</li> </ul> | <ul style="list-style-type: none"> <li>• Computer</li> <li>• Group creator (by group clauses only)</li> <li>• Group name (by group clauses only)</li> <li>• Instance</li> <li>• Database-tablespace</li> <li>• User name</li> <li>• Space used</li> <li>• Threshold</li> <li>• % of threshold</li> <li>• Number of violations</li> <li>• Quota creator</li> <li>• Quota name</li> <li>• Quota type</li> </ul> |

## Data Manager for Chargeback reports

Use invoices generated by Data Manager for Chargeback to charge users or groups of users for storage usage within a storage environment.

See “Working with invoices” on page 625 for information about how to use Data Manager for Chargeback to generate invoices.

## Disk Manager reports

Use Disk Manager reports to view the relationships among the components of a storage subsystem and the storage capacity for a computer, filesystem, storage subsystem, volume, Array Sites, Ranks, and disk.

### Reporting on storage subsystems

Learn how to use Disk Manager to report on storage subsystems.

These reports also enable you to view the relationships among the components of a storage subsystem. Storage-subsystems reports use data gathered by probes.

To ensure the accuracy and consistency of the data that appears in storage-subsystem reports, run regularly scheduled probes against the hosts that use or import volumes, and the storage subsystems upon which you want to report. You should run probes against storage subsystems after running them against any host.

- A disk must be partitioned before it is considered a logical volume by Disk Manager. If a disk is not partitioned, it will not be included in a file system or logical volumes report.

- Before you begin, you might want to refer to Generating and viewing reports. That section provides information on reports that applies all IBM Tivoli Storage Productivity Center reports. It describes the types of reports: system, user-defined, and batch. It also describes how to save reports and how to customize reports, such as filtering information and displaying the information in chart form.

You can view storage subsystem reports in the following ways:

**Disks** Use these reports to view and chart the storage capacity of storage subsystem disks. Use **By** clauses to filter the storage entities that appear in a report. For example, if you want to view a report that displays specific disks, select **By Disk**; if you want to view the disks of specific storage subsystems, select **By Storage Subsystems**.

#### **Volumes**

Use these reports to view and chart the storage capacity of storage subsystem volumes. Use the **By Storage Subsystem** and **By Storage Subsystem Group** clauses to filter the storage entities that appear in a report; use the **By Volume Space** and **By Probe Time** clauses to sort the volumes that appear in a report according to when a volume was probed or the capacity of a volume.

#### **Storage Pools**

Use these reports to view and chart the storage capacity of storage pools. Use **By** clauses to filter the storage entities that appear in a report. For example, if you want to view a report that displays information about the storage pools for specific storage systems, select **By Storage Subsystem**.

#### **Array Sites**

Use these reports to view information about the array sites that reside on storage subsystems. Use **By** clauses to filter the storage entities that appear in a report. For example, if you want to view a report that displays about the array sites for specific groups of storage subsystems, select **By Storage Subsystem Group**.

#### **Disk Groups**

Use these reports to view information about the disk groups that reside on storage subsystems. Use the **By** clauses to filter the storage entities that appear in a report. For example, if you want to view a report that displays about the disk groups for specific storage subsystems, select **By Storage Subsystem**.

**Ranks** Use these reports to view information about the ranks that reside on storage subsystems. Use the **By** clauses to filter the storage entities that appear in a report. For example, if you want to view a report that displays about the ranks for specific groups of storage subsystems, select **By Storage Subsystem Group**.

#### **Computer Views**

Use these reports to view the relationships between the computer and the storage subsystems, volumes, and disks in your environment. The information is gathered from the computer's viewpoint. Note that, for reporting purposes, *computer* is defined as anything that is running Tivoli Storage Productivity Center.

**Note:** These reports are available only if the following steps have been performed:



- A Storage Resource agent or Data agent is installed on a machine where a volume is surfaced
- A probe has been run against the machine where the volume is located
- A probe has been run against the storage subsystem or IBM SAN Volume Controller cluster

### By Computer

These reports view relationships between the computer and the storage subsystems, volumes, or disks in your environment. To generate a report by computer, select the **By Computer** node and choose an item from the **Relate Computers to:** field.

For example, to retrieve information on computers in your environment and their volumes, select the **By Computer** node and then choose **By volumes**. These settings will generate a report detailing all computers with agents on them, which use volumes from the monitored storage subsystems.

### By Computer Group

These reports view relationships between the computer groups (both monitoring and reporting groups) and the storage subsystems, volumes, or disks in your environment. To generate a report by computer group, select the **By Computer group** node and choose an item from the **Relate Computers to:** field.

For example, to retrieve information on computers in a group that use storage from monitored storage subsystems, select the **By Computer Group** node and then choose **Storage Subsystems** from the **Relate Computers to:** field. These settings will generate a report detailing all computers in groups with agents on them, which use volumes from the monitored storage subsystems.

### By Filesystem/Logical Volume

These reports will view the relationships between the filesystems or logical volumes and the storage subsystems, volumes, and disks in your environment. To generate a report for the relationship you want to view, select the **By Filesystem/Logical Volume** node and choose an item from the **Relate Filesystem/Logical Volume to:** field.

For example, to retrieve information on filesystems or logical volumes that use storage from monitored storage subsystems, select the **By Filesystem/Logical Volume** node and choose **By Storage Subsystems**. These settings will generate a report detailing all filesystems or logical volumes which use storage from monitored storage subsystems.

### By Filesystem Group

Use this report to view the relationships between the filesystem groups (both monitoring and reporting) and the storage subsystems, volumes, and disks in your environment. To generate a report for the relationship you want to view, select the **By Filesystem Group** node and choose an item from the **Relate Filesystem/Logical Volume to:** field.

For example, to retrieve information on filesystems in a group that use storage from monitored storage subsystems, select the **By Filesystem Group** node and choose **Storage Subsystems** from the **Relate Filesystem/Logical Volume to:** field. These settings will

generate a report detailing all filesystems in a group which use storage from monitored storage subsystems.

### Storage Subsystem Views

Use these reports to view the relationships among the storage-subsystem components (storage subsystems, volumes, and disks) to the computers and file systems or logical volumes in your environment.

**Note:** These reports are available only if the following steps have been performed:

- A Data agent or Storage Resource agent is installed on a system where a volume is surfaced
- A probe has been run against the system where the volume is located
- A probe has been run against the storage subsystem or IBM SAN Volume Controller cluster
- For IBM XIV Storage System, there is no relationship between storage disks, pools, and volumes because of the unique storage disk distribution and architecture.

#### By Storage Subsystem

These reports view the relationships between the storage subsystems and the computers and file systems or logical volumes in your environment. To generate a report for the relationship you want to view, select the **By Storage Subsystem** node and choose an item from the **Relate Storage Subsystem to:** field.

For example, to view all storage subsystems which have storage being used by computers in the environment, select **By Storage Subsystem** and choose **Computers**. These settings will generate a report detailing all storage subsystems which have storage being used by computers in the environment.

#### By Storage Subsystem Group

These reports view the relationships between storage subsystem groups (both monitoring and reporting) and the computers and file systems or logical volumes in your environment. To generate a report for the relationship you want to view, select the **By Storage Subsystem Group** node and choose an item from the **Relate Storage Subsystem to:** field.

For example, to view all storage subsystem groups that have storage being used by computers in the environment, select **By Storage Subsystem Group** and choose **Computers** from the **Relate Storage Subsystem to:** field. These settings will generate a report detailing all storage subsystem groups that have storage being used by computers in the environment.

#### By Volume

Use this report to view the relationships between the volumes and the computers and file systems or logical volumes in your environment. To generate a report for the relationship you want to view, select the **By Volume** node and then choose an item from the **Relate Volumes to:** field.

For example, to get a listing of all volumes in the environment on which file systems or logical volumes have been created, select **By Volume** and choose **Filesystems/Logical Volumes** from the **Relate**

**Volumes to:** field. These settings will generate a report detailing all volumes in the environment on which file systems or logical volumes have been created.

To get a listing of all volumes in the environment which are being used by computers in the environment, select **By Volume** and choose **Computers**. These settings will generate a report detailing all volumes in the environment which are being used by computers in the environment.

**Disks** Use this report to view the relationships between disks and the computers and file systems or logical volumes in your environment. To generate a report for the relationship you want to view, select the **Disks** node and then choose an item from the **Relate Disks to:** field.

For example, to get a listing of all storage subsystem disks in the environment on which file systems or logical volumes have been created, select **Disks** and choose **Filesystems/Logical Volumes**. These settings will generate a report detailing all storage subsystem disks in the environment on which file systems or logical volumes have been created.

### **Volume to HBA Assignment**

Use these reports to view all the logical unit numbers (volumes) of each probed storage subsystem that has been assigned to at least one host bus adaptor (HBA) port world wide name (WWN). These reports only retrieve data from machines that have HBAs.

The HBA is a fibre channel and volumes are physically attached to the storage subsystem.

#### **By Storage Subsystem**

This report shows all the volumes of each probed storage subsystem that has been assigned to at least one HBA Port WWN. This report is sorted by storage subsystem name.

#### **By Volume Space**

This report views all the volumes of each probed storage subsystem that has been assigned to at least one HBA Port WWN. This report is sorted by volume space.

#### **Not on Monitored Computers**

This report shows only volumes that have been assigned to HBA Port WWNs on hosts with no Storage Resource agent or Data agent installed on them.

#### **Not Visible to Monitored Server**

Use this report to find volumes that have been assigned to hosts with Storage Resource agents or Data agents, but not configured. This report can be used to generate a list of all volumes that need to be configured before they are used.

### **Volume to Backend Volume Assignment By Volume**

Use this report to view the relationships between the virtual disks (including mirror copies) and managed disks that are associated with a virtualizer storage subsystem, such as the SAN Volume Controller, and the back-end storage volumes that are associated with those managed disks.

**Note:** A virtual disk with one mirror copy appears as one row in the table on this page. A virtual disk with multiple mirror copies appears in multiple rows in the table. Each row represents one of the copies.

**Important:**

Some of the columns in the totals row for these reports can count storage multiple times if the report includes the following:

- A storage virtualizer like a SAN Volume Controller
- The disk arrays that host the backend storage to a storage virtualizer like a SAN Volume Controller

In this situation, disks and logical unit numbers (LUNs) appear to belong to both a storage virtualizer like SAN Volume Controller and the disk arrays, and are counted multiple times in some of the columns in the totals row. To avoid this situation, filter the report to only display the storage virtualizer or the disk arrays. If this situation is unavoidable, then some of the columns in the totals row can be adjusted as follows:

1. Use the **Physical Disk Space** column to determine the actual physical disk space on all arrays. This disk space is not counted multiple times.
2. Sum the **Disk Space** column for the disk arrays and subtract this result from the totals row to determine the total backend storage for a virtualizer.
3. To prevent counting the backend volumes multiple times, subtract the **Backend Volume Space** column from the **Volume Space** column in the totals row to determine the amount of volume space.

## Generating storage-subsystem reports



Learn how to generate a storage subsystem report.

These reports are available only if the following steps have been performed:

1. An agent is installed on a machine where a volume is surfaced.
2. A probe has been run against the machine where the volume is located.
3. A probe has been run against the storage subsystem or SAN Volume Controller cluster.

To generate a storage-subsystems report, do the following:

1. In the left pane, expand **Disk Manager > Reporting > Storage Subsystems**.
2. Do one of the following:
  - a. In the left pane, expand **Computer Views** and click one of the following:
    - By Computer
    - By Computer Group
    - By Filesystem/Logical Volume
    - By Filesystem Group
  - b. In the left pane, expand **Storage Subsystems Views** and click one of the following:
    - By Storage Subsystem
    - By Storage Subsystem Group
    - By Volume
    - Disks
  - c. In the left pane, expand **Volume to HBA Assignment** and click one of the following:

- By Storage System
  - By Volume Space
  - Not On Monitored Computers
  - Not Visible to Monitored Computers
- d. In the left pane, expand **Volume to Backend Volume Assignment** and click the following:
- By Volume
3. The **Storage Subsystem Performance** pane for your selected storage subsystem opens the Select Resources pane. All valid available resources are listed in Included Columns. In Available Columns and Included Columns, use  and  to select the columns you want to include in your report.
4. Filters can be applied to the selected columns. In the **Report Filter Specification** area, the **Filter** button opens the **Edit Filter** window.
- a. Click **Add** to add a column to the filter list. The columns will be added in the order they appear on your report.
  - b. Click **Delete** to delete all columns. Click the X to the left of an individual column to delete it.
  - c. In the **Records Must Meet** area, select **All Conditions** or **At Least One Condition**.
  - d. Type the condition value in the **Value 1** text field.
    - If the text field is greyed-out, click **Edit**.
    - If you wish to add additional values for the same column using the same operator, click **Edit**.
  - e. Select the Operator.
    - If you select the BETWEEN or NOT BETWEEN operators, two value text fields will open for a lower and upper limit. Click **Edit** to enter the upper and lower limit. You must input values in both **Value 1** and **Value 2** fields, if shown.
    - For numeric columns, such as Free Space, a pull-down menu at the end of the each field assists you with numerical units.
  - f. Once your selections have been made, click **OK**.
5. Click **Generate Report**. The report appears under a new tab window. If there is no data, as in if a probe has not been run, the window will say "No Data To Display."
6. Click **File > Save**. The report is saved under **IBM Tivoli Storage Productivity Center > My Reports**. Reports can be regenerated from that node.

To change the parameters for your report, or create a new report of the same type, do the following:

1. Click the **Selection** tab before closing the Storage Subsystem Performance window. Make the necessary changes.
2. Click **Generate Report**.
3. Click **File > Save** and enter the name of the new report.
4. Click **OK**.

## Creating a storage subsystem performance report

Learn how to create a storage-subsystem performance report.

Before you can view a performance report on a storage subsystem, that subsystem must have been discovered and had performance data successfully collected. For more information on storage-subsystem discovery, see "Discover storage subsystem, tape library, and fabric information with CIM agents" on page 330.

For a complete list of metrics available for performance reports, see "Performance metrics" on page 1186.

1. In the left pane, expand **Disk Manager > Reporting > Storage Subsystem Performance**.
2. Click the type of performance report that you want to view. You can choose from the following performance reports:

**By Storage Subsystem**

This report contains performance information for storage subsystems.

**By Controller**

This report contains performance information for controllers.

**By I/O Group**

This report contains performance information for I/O groups.

**By Node**

This report contains historical performance data for SAN Volume Controller (SVC) nodes.

**By Array**

This report contains performance information for arrays.

**By Managed Disk Group**

This report contains performance information for managed-disk groups.

**By Volume**

This report contains performance information for volumes.

**By Managed Disk**

This report contains performance information for managed disks.

**By Port**

This report contains performance information for ports.

In the right pane, the Selection page opens.

3. Select from the following options to customize the report:

**Selection button**

Click this button to determine the resources upon which you want to report, such as the storage subsystems, the fibre-channel ports, and so forth. The Select Resources window displays. The type of resources displayed for selection depends on the type of report you selected. For example, if you selected a By Storage Subsystem report, all displayed resources will be storage subsystems. If you clear the box for a resource on the Select Resources window, a pushpin icon displays on the Selection button.

**Filter button**

Click this button to further defines the data that you want to include in the report. Filters enable you to apply general rules to the report. When you click this button, the Edit Filter window displays.

For example, to select only data associated with subsystems that have a name ending with "123-IBM", select the "Subsystem" column, the "LIKE" operator, and type "\*123-IBM" (without the quotation marks) in

the Value 1 field. To select only data which has a Read I/O Rate greater than 0, select the "Read I/O Rate (overall)" column, the ">" operator, and type "0" (without the quotation marks) in the Value 1 field.

If a filter is defined for this report, a pushpin icon displays on the **Filter** button.

### Time range selectors

All performance reports display historical performance data. Use the time range selectors to specify the exact time range that you want the report to cover. Click one of the following radio buttons to select the type of historical report you want:

#### Display latest performance data

This report contains only the latest available performance data for the selected resources. Note that even the latest data available might still be somewhat old, depending on whether data was recently collected for the resources. Including the time column in the report will help you determine how old the latest available data actually is. Regardless of the age of the data, clicking this button ensures that only a single row (representing a single point in time) for each selected resource is included in the report. If no data has ever been collected for a particular resource, then that resource is not included in the report. If no data has ever been collected for any resource, then the resulting report will be empty.

#### Display historic performance data using absolute time

This report contains only data that falls within a specific time range. Select the time range using the **From** and **To** selectors. The specified from and to values are interpreted as inclusive. The resulting report contains all the performance data that is available within the specified time range, for each selected resource. This means that the data can consist of potentially many rows per resource, each with a unique timestamp. If no data has been collected for a particular resource during the specified time range, that resource is not included in the report. If no data has been collected for any resource during the specified time range, then the resulting report will be empty.

To determine which historical performance records fall within the specified time range, the **beginning timestamp** of the collection intervals is used for comparison. If an historical record contains performance data for an interval of time from 1:00 PM to 1:05 PM, that record would be included in the report for a time range of 1:00 PM to 8:00 PM. However, that same record would not be included if a time range of 1:02 PM to 8:02 PM was specified, because it is the beginning of the time period (for example, 1:00 PM in this example) that is used to determine inclusion in the specified time interval.

#### Display historic performance data using relative time

This report contains only data that falls within a specific time range extending backward from the current date and time. Specify the number of days in the desired time range, using the **days ago until now** field. The report will be exactly the same as a report using absolute time, with the **To** selector set to the current date and time, and the **From** selector set to the current date and time minus the number of **days ago until now**.

Note that the same criteria applies as if determining which historical performance records fall within the specified time range. Only the beginning timestamp of the collection intervals is used for comparison. This means that data for the current hour or the current day is not included in the report when selecting hourly or daily summation data, because data for the current hour or current day is only saved in the database when the hour or day is complete. As an alternative, consider using the first radio button, **Display latest performance data**, to see a report showing the most recent daily summation data.

### Summation Level

Use this selector to select which type of performance data records to include in the report. The Tivoli Storage Productivity Center Performance Manager collects and stores three different types of performance data:

#### By sample

Represents the most detailed data. There is usually one sample record per resource for every **n** minutes, where **n** is the performance monitor interval length that was specified when the monitor was defined. Each sample represents the average performance of the resource over those **n** minutes. A sample record with a timestamp of **t** will be saved at time **t+n**. In other words, the beginning of each **n** minute time interval is used for display and for time range matching.

#### Hourly

Represents an hourly average of the performance of the specified resources. There will be one hourly record written every hour, per resource. Each such record will be written at or shortly after each hour mark on the (device) clock. For example, a record with a timestamp of 12:00 PM and interval length of 1 hour will be written at 1:00 PM.

**Daily** Represents a daily average of the performance of the resources. There will be one daily record written every day, per resource. Each daily record will be written at or shortly after midnight, according to the device clock. For example, a record with a timestamp of April 21, 12:00 AM and interval length of 24 hours will be written on April 22, 12:00 AM.

Note that hourly and daily records are usually retained in the database for longer periods of time, compared to By Sample records. As a result, if a generated report indicates that sample data is no longer available for a particular time range in the past, it is possible that the sample data has already been purged, in which case you can try generating the report using hourly or daily summation levels.

Be aware that when selecting a time range using the **Display historic performance data using relative time** radio button, it is necessary to select the length of the range in days. This can lead to unexpected results when also selecting the daily summation level. Remember that a daily record is not written until the day is complete, which means that at any given time during the day, the current day's data does not yet exist in the database. Therefore, selecting the daily summation level as well as a relative time range of **n days ago until now**, will usually result in **n-1** records being returned per resource, assuming that at least



n-1 days worth of data has been previously collected for those resources. Specifying a relative time range of 1 **days ago until now** will result in an empty report.

**Note:** The user interface displays the date and time of data in the Tivoli Storage Productivity Center server time zone, but the daily and hourly aggregations are calculated based on the device time zones. Hourly records are written at the end of each hour according to the devices' clocks and the daily records are written at midnight according to the devices' clocks, but both are displayed according to the Tivoli Storage Productivity Center server's clock. For example, if a subsystem clock is one hour ahead of the Tivoli Storage Productivity Center server clock, the hourly data for 5:00 PM according to the device's clock is displayed as 4:00 PM in the user interface and the daily data for 12:00 AM according to the device's clock is displayed as 11:00 PM of the previous day.

### **Available Columns, Included Columns**

Use these items to specify which columns to include and which columns to exclude from the generated report.

#### **Available Columns**

Displays the columns that you can include in the generated report that are not already included. If a column appears in the **Included Columns** list box, it will not appear in the **Available Columns** list box.

#### **Included Columns**

Displays the columns that will appear in the generated report. By default, all columns for a report are listed in this list box and will be displayed in the generated report. You can highlight the columns and use the up and down arrows to rearrange the order of the columns in the report. The order of the items in this list box determines the order in which the columns will appear in the generated report. The grayed-out column names shown at the top of the list box are the frozen columns, which cannot be selected, removed, or reordered for this type of report.



If a column appears in the **Available Columns** list box, it will not appear in the **Included Columns** list box.

To exclude a column from a report, click the name of the column in the **Included Columns** list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the left arrow button (<<). The selected columns are removed from the **Included Columns** list box, and will appear in the **Available Columns** list box.

To include a column in a report, click the name of the column in the **Available Columns** list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the right arrow button (>>). The selected columns will be removed from the **Available Columns** list box, and will appear in the **Included Columns** list box.

To rearrange the order of columns that will appear in the generated report, click the name of a column in the **Available Columns** list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so

desired. Then click the up arrow button (^) or the down arrow button (v) to move the selected columns up or down in the list, relative to the other columns.

4. Click **Generate Report**. A new tab opens and displays the report.
5. Optional: To view the report in chart format, click , specify charting options, and then click **OK**. Right-click the chart and select **Customize this Chart** to further customize how the chart is displayed.
6. Optional: To save the report settings to generate later, click  and enter a name for the report. The report is saved under the IBM Tivoli Storage Productivity Center node. See “Filtering the resources that appear on a report” on page 799 for instructions on modifying saved report settings.




If the report is greater than 2500 rows, a dialog box will appear which will allow you to retrieve additional rows by clicking **Display More Rows**. An additional 2500 rows will be retrieved from the server and displayed on the panel for each click of the button. Do not try to display too much data or there is the possibility of running out of memory.

## Generating constraint violation reports

Learn how to generate a storage subsystem performance report.

### Creating constraint violation reports:

Learn how to how to create constraint violation reports.

1. In the left pane, expand **Disk Manager > Reporting > Storage Subsystem Performance**.
2. Click **Constraint Violations**. The constraint violation report selection page appears in the right pane.
3. If you want to limit the report to a particular time period, select **Limit to this time range** and specify the time period.
4. Click **Selection** to choose which storage subsystems to include in the report.
5. Click **Filter** to modify the filter used in generating the report.
6. Customize the performance report. By default, all supported metrics are included. Click to move the selected metric up in the list and to move the selected metric down in the list.
  - a. In the **Included Columns** list, select a metric and click  to remove it from the report.
  - b. To change the order in which the columns appear in the report, select a metric and click either  or .
7. Click **Generate Report**. The **Constraint Violations** tab opens and displays the report.

The report might show information about DS6000, DS8000, TotalStorage Enterprise Storage Server, SAN Volume Controller, and other subsystems all on the same report. By default, the report displays one subsystem per row. Results are aggregated for the current day.


### Viewing detailed constraint violation information:

Learn how to view detailed information about constraint violations.

Before you can view detailed information about constraint violations, you must generate a constraint violation report (see “Creating constraint violation reports” on page 972).

1. Click the **Constraint Violations** tab to view the report. The report might show information about DS6000, DS8000, TotalStorage Enterprise Storage Server, SAN Volume Controller, and other subsystems all on the same report. By default, the report displays one subsystem per row. Results are aggregated for the current day.


2. Click  to view a detailed report of that subsystem.


3. Click , specify charting options, and then click **OK** to view the report in chart form. Right-click the chart and select **Customize this Chart** to further customize how the chart is displayed.

### Viewing affected volumes and affected hosts:

Learn how to view information on volumes and related hosts affected by constraint violations.

Before you can view affected volumes and hosts information, you must generate a constraint violation report (see “Creating constraint violation reports” on page 972) and view detailed information on a subsystem (see “Viewing detailed constraint violation information” on page 972).

1. Click  beside a constraint violation to display the Affected Volumes Selection page.
2. Under Volumes, select whether you want the report to show all volumes or only the most active volumes associated with the subsystem and component and type in the maximum number of volumes to display in the report.
3. Under Performance Data, select whether you do not want the report to display performance data for the volumes or you do want the report to display historic performance data for volumes.
4. Customize the performance report. By default, all supported metrics are included. Click to move the selected metric up in the list and to move the selected metric down in the list.

- a. In the **Included Columns** list, select a metric and click  to remove it from the report.

- b. To change the order in which the columns appear in the report, select a

metric and click either  or .

5. Click **Generate Report**.

## Fabric Manager reports

Fabric reports provide you with both overview and detailed information about your fabrics, switches, and ports. The reports are based on information collected by probes and performance monitors. You can view reports as tables of data or as graphical charts. You can use pre-defined Fabric system reports or define your own switch performance reports for information by port or by constraint violations.

Before you can generate a report you need to gather information about your fabrics. You can do this through monitoring jobs, and discoveries.

## Creating switch performance reports

After performance monitors are in operation, you can generate switch performance reports. The procedure is similar for each report.

“Creating a switch performance monitor” on page 375 describes how to create performance monitors for fabric switches.

For both reports, use the **Selection** tab to specify switches to report, report columns to include, and filters to apply. When you click **Generate Report**, the **Ports** page or **Constraint Violations** page is created and displays the report.

1. Expand **Fabric Manager** → **Reporting** → **Switch Performance**.
2. Click **By Switch**, **By Port** or **By Constraint Violations**. In the right pane, the Selection page opens.
3. Select from the following options to customize the report:

### Selection button

Click this button to determine the resources upon which you want to report. The Select Resources window displays. The type of resources displayed for selection depends on the type of report you selected. If you clear the box for a resource on the Select Resources window, a pushpin icon displays on the Selection button.

### Filter button

Click this button to further defines the data that you want to include in the report. Filters enable you to apply general rules to the report. When you click this button, the Edit Filter window displays.

For example, to select only data associated with switches that have a name ending with "123-IBM", select the "Switch" column, the "LIKE" operator, and type "\*123-IBM" (without the quotation marks) in the **Value 1** field.

If a filter is defined for this report, a pushpin icon displays on the **Filter** button.

### Time range selectors

All performance reports display historical performance data. Use the time range selectors to specify the exact time range that you want the report to cover. Click one of the following radio buttons to select the type of historical report you want:

#### Display latest performance data

This report contains only the latest available performance data for the selected resources. Note that even the latest data available might still be somewhat old, depending on whether data was recently collected for the resources. Including the time column in the report will help you determine how old the latest available data actually is. Regardless of the age of the data, clicking this button ensures that only a single row (representing a single point in time) for each selected resource is included in the report. If no data has ever been collected for a particular resource, then that resource is not included in the report. If no data has ever been collected for any resource, then the resulting report will be empty.

### Display historic performance data using absolute time

This report contains only data that falls within a specific time range. Select the time range using the **From** and **To** selectors. The specified from and to values are interpreted as inclusive. The resulting report contains all the performance data that is available within the specified time range, for each selected resource. This means that the data can consist of potentially many rows per resource, each with a unique timestamp. If no data has been collected for a particular resource during the specified time range, that resource is not included in the report. If no data has been collected for any resource during the specified time range, then the resulting report will be empty.

To determine which historical performance records fall within the specified time range, the **beginning timestamp** of the collection intervals is used for comparison. If an historical record contains performance data for an interval of time from 1:00 PM to 1:05 PM, that record would be included in the report for a time range of 1:00 PM to 8:00 PM. However, that same record would not be included if a time range of 1:02 PM to 8:02 PM was specified, because it is the beginning of the time period (for example, 1:00 PM in this example) that is used to determine inclusion in the specified time interval.

### Display historic performance data using relative time

This report contains only data that falls within a specific time range extending backward from the current date and time. Specify the number of days in the desired time range, using the **days ago until now** field. The report will be exactly the same as a report using absolute time, with the **To** selector set to the current date and time, and the **From** selector set to the current date and time minus the number of **days ago until now**.

Note that the same criteria applies as if determining which historical performance records fall within the specified time range. Only the beginning timestamp of the collection intervals is used for comparison. This means that data for the current hour or the current day is not included in the report when selecting hourly or daily summation data, because data for the current hour or current day is only saved in the database when the hour or day is complete. As an alternative, consider using the first radio button, **Display latest performance data**, to see a report showing the most recent daily summation data.

### Summation Level

Use this selector to select which type of performance data records to include in the report. The Tivoli Storage Productivity Center Performance Manager collects and stores three different types of performance data:

#### By sample

Represents the most detailed data. There is usually one sample record per resource for every **n** minutes, where **n** is the performance monitor interval length that was specified when the monitor was defined. Each sample represents the average performance of the resource over those **n** minutes. A sample record with a timestamp of **t** will be saved at time **t+n**. In other

words, the beginning of each **n** minute time interval is used for display and for time range matching.

### **Hourly**

Represents an hourly average of the performance of the specified resources. There will be one hourly record written every hour, per resource. Each such record will be written at or shortly after each hour mark on the (device) clock. For example, a record with a timestamp of 12:00 PM and interval length of 1 hour will be written at 1:00 PM.

**Daily** Represents a daily average of the performance of the resources. There will be one daily record written every day, per resource. Each daily record will be written at or shortly after midnight, according to the device clock. For example, a record with a timestamp of April 21, 12:00 AM and interval length of 24 hours will be written on April 22, 12:00 AM.

Note that hourly and daily records are usually retained in the database for longer periods of time, compared to By Sample records. As a result, if a generated report indicates that sample data is no longer available for a particular time range in the past, it is possible that the sample data has already been purged, in which case you can try generating the report using hourly or daily summation levels.

Be aware that when selecting a time range using the **Display historic performance data using relative time** radio button, it is necessary to select the length of the range in days. This can lead to unexpected results when also selecting the daily summation level. Remember that a daily record is not written until the day is complete, which means that at any given time during the day, the current day's data does not yet exist in the database. Therefore, selecting the daily summation level as well as a relative time range of **n days ago until now**, will usually result in **n-1** records being returned per resource, assuming that at least **n-1** days worth of data has been previously collected for those resources. Specifying a relative time range of **1 days ago until now** will result in an empty report.

**Note:** The user interface displays the date and time of data in the Tivoli Storage Productivity Center server time zone, but the daily and hourly aggregations are calculated based on the device time zones. Hourly records are written at the end of each hour according to the devices' clocks and the daily records are written at midnight according to the devices' clocks, but both are displayed according to the Tivoli Storage Productivity Center server's clock. For example, if a subsystem clock is one hour ahead of the Tivoli Storage Productivity Center server clock, the hourly data for 5:00 PM according to the device's clock is displayed as 4:00 PM in the user interface and the daily data for 12:00 AM according to the device's clock is displayed as 11:00 PM of the previous day.

### **Available Columns, Included Columns**

Use these items to specify which columns to include and which columns to exclude from the generated report.

#### **Available Columns**

Displays the columns that you can include in the generated

report that are not already included. If a column appears in the **Included Columns** list box, it will not appear in the **Available Columns** list box.

### Included Columns

Displays the columns that will appear in the generated report. By default, all columns for a report are listed in this list box and will be displayed in the generated report. You can highlight the columns and use the up and down arrows to rearrange the order of the columns in the report. The order of the items in this list box determines the order in which the columns will appear in the generated report. The grayed-out column names shown at the top of the list box are the frozen columns, which cannot be selected, removed, or reordered for this type of report.

If a column appears in the **Available Columns** list box, it will not appear in the **Included Columns** list box.


To exclude a column from a report, click the name of the column in the **Included Columns** list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the left arrow button (<<). The selected columns are removed from the **Included Columns** list box, and will appear in the **Available Columns** list box.

To include a column in a report, click the name of the column in the **Available Columns** list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the right arrow button (>>). The selected columns will be removed from the **Available Columns** list box, and will appear in the **Included Columns** list box.

To rearrange the order of columns that will appear in the generated report, click the name of a column in the **Available Columns** list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the up arrow button (^) or the down arrow button (v) to move the selected columns up or down in the list, relative to the other columns.

4. Click **Generate** to run the report.

The report is displayed in the **Ports** tab. You can display the information in

chart form by clicking .

5. Click **File** → **Save as** to save the report in the <user-id> **Reports** node.

### Switch Performance:

Use the Switch Performance report to generate and view reports that provide information about the free and occupied ports on the switches in your SAN. These reports provide information about the port connections and status, as well as the port and switch IDs, data and packet transfer rates, and other details. This report can only be generated for switches that have had performance monitors run on them.

This following table describes the columns on the Switch Performance report. All the columns in the report are for display only.


---

|        |                       |
|--------|-----------------------|
| Switch | The ID of the switch. |
|--------|-----------------------|

---

|                                    |                                                                            |
|------------------------------------|----------------------------------------------------------------------------|
| <b>Time</b>                        | The date and time of the data collection.                                  |
| <b>Interval</b>                    | The size of the sample interval, in seconds.                               |
| <b>Port Send Packet Rate</b>       | The rate at which packets are sent through this port.                      |
| <b>Port Receive Packet Rate</b>    | The rate at which packets are received through this port.                  |
| <b>Total Port Packet Rate</b>      | The total rate at which packet operations are performed through this port. |
| <b>Port Send Data Rate</b>         | The rate at which data is sent through this port.                          |
| <b>Port Receive Data Rate</b>      | The rate at which data is received through this port.                      |
| <b>Total Port Data Rate</b>        | The total rate at which data is processed through this port.               |
| <b>Port Peak Send Data Rate</b>    | The peak rate at which data is sent through this switch.                   |
| <b>Port Peak Receive Data Rate</b> | The peak rate at which data is received through this switch.               |
| <b>Port Send Packet Size</b>       | The size of the data block that is sent out through this switch.           |
| <b>Port Receive Packet Size</b>    | The size of the data block that is received in through this switch.        |
| <b>Overall Port Packet Size</b>    | The total size of the data block that is transferred by the switch.        |
| <b>Error Frame Rate</b>            | The rate at which error frames are generated by the switch.                |
| <b>Dumped Frame Rate</b>           | The rate at which frames are dumped by the switch.                         |
| <b>Link Failure Rate</b>           | The rate at which failed links are generated by the switch.                |

### Fabric Performance Report - By Port:

Click  to display detailed information per computer in a **Network** tab. The **Network** tab displays a file summary by computer.

The following fields on the Fabric Performance Report - By Port are for display only.

|                                    |                                                                                                                                                                                                                    |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Port</b>                        | The ID of the port.                                                                                                                                                                                                |
| <b>WWPN</b>                        | The World Wide Port Name (WWPN) of the selected port.                                                                                                                                                              |
| <b>Slot</b>                        | The number for the blade that the port is on.                                                                                                                                                                      |
| <b>Index</b>                       | The unique number for the port across the whole switch.                                                                                                                                                            |
| <b>Time</b>                        | The date and time of the data collection.                                                                                                                                                                          |
| <b>Interval</b>                    | The size of the sample interval, in seconds.                                                                                                                                                                       |
| <b>Port Send Packet Rate</b>       | The average number of packets sent per second, by a particular port over a particular time interval. A send operation is a read operation processed, or a write operation initiated by the particular port.        |
| <b>Port Receive Packet Rate</b>    | The average number of packets received per second, by a particular port over a particular time interval. A receive operation is a write operation processed, or a read operation initiated by the particular port. |
| <b>Total Port Packet Rate</b>      | The average number of packets sent and received per second, by a particular port over a particular time interval.                                                                                                  |
| <b>Port Send Data Rate</b>         | The average number of megabytes sent per second, by a particular port over a particular time interval.                                                                                                             |
| <b>Port Receive Data Rate</b>      | The average number of megabytes received per second, by a particular port over a particular time interval.                                                                                                         |
| <b>Total Port Data Rate</b>        | The average number of megabytes transferred per second, by a particular port over a particular time interval.                                                                                                      |
| <b>Port Peak Send Data Rate</b>    | The peak number of megabytes sent per second, by a particular port over a particular time interval.                                                                                                                |
| <b>Port Peak Receive Data Rate</b> | The peak number of megabytes received per second, by a particular port over a particular time interval.                                                                                                            |
| <b>Port Send Packet Size</b>       | The size of the data block that is sent out through this switch.                                                                                                                                                   |
| <b>Port Receive Packet Size</b>    | The size of the data block that is received in through this switch.                                                                                                                                                |
| <b>Overall Port Packet Size</b>    | The total size of the data block that is transferred by the switch.                                                                                                                                                |
| <b>Error Frame Rate</b>            | The rate at which error frames are generated by the switch.                                                                                                                                                        |




|                                      |                                                                                                                                                                                                                                      |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Dumped Frame Rate</b>             | The rate at which frames are dumped by the switch.                                                                                                                                                                                   |
| <b>Link Failure Rate</b>             | The rate at which failed links are generated by the switch.                                                                                                                                                                          |
| <b>Loss of Sync Rate</b>             | The average number of times per second that synchronization was lost after the last reset of the device, for a particular component over a particular time interval.                                                                 |
| <b>Loss of Signal Rate</b>           | The average number of times per second that the signal was lost after the last reset of the device, for a particular component over a particular time interval.                                                                      |
| <b>CRC Error Rate</b>                | The average number of frames received per second in which the CRC in the frame did not match the CRC computed by the receiver, for a particular component over a particular time interval.                                           |
| <b>Short Frame Rate</b>              | The average number of frames received per second that were shorter than 28 octets (24 header + 4 CRC), not including any start of frame/end of frame bytes, for a particular component over a particular time interval.              |
| <b>Long Frame Rate</b>               | The average number of frames received per second that were longer than 2140 octets (24 header + 4 CRC + 2112 data), not including any start of frame/end of frame bytes, for a particular component over a particular time interval. |
| <b>Encoding Disparity Error Rate</b> | The average number of disparity errors received per second, for a particular component over a particular time interval.                                                                                                              |
| <b>Discarded Class3 Frame Rate</b>   | The average number of class-3 frames per second that were discarded by a particular component over a particular time interval.                                                                                                       |
| <b>F-BSY Frame Rate</b>              | The average number of F-BSY frames per second that were generated by a particular component over a particular time interval.                                                                                                         |
| <b>F-RJT Frame Rate</b>              | The average number of F-RJT frames per second that were generated by a particular component over a particular time interval.                                                                                                         |

### Constraint Violations:

If you have defined switch performance alerts you can also generate the Constraint Violations report.



Click  to display detailed information per computer in a **Network** tab. The **Network** tab displays a file summary by computer.

The following fields on the Top Switch Ports Packet Rate Performance report are for display only (not editable).

|                                         |                                                                                |
|-----------------------------------------|--------------------------------------------------------------------------------|
| <b>Switch</b>                           | The ID of the switch.                                                          |
| <b>Total Port Data Rate Threshold</b>   | The total rate at which data is processed through this port.                   |
| <b>Link Failure Rate Threshold</b>      | The rate at which failed links are generated by the switch.                    |
| <b>Error Frame Rate Threshold</b>       | The rate at which error frames are generated by the switch.                    |
| <b>Total Port Packet Rate Threshold</b> | The total rate at which the packet operations are performed through this port. |

## Tape Manager reports

You can perform the following tape library reporting tasks:

- Create and view asset reports using data collected from your tape libraries; these reports can include such information as the library identifier and the status of the library
- Create and view capacity reports using data collected from your tape libraries; these reports can include such information as the number of drives in a library, the number of tape slots, the number of occupied tape slots

---

## System reports

Use the information in this section to learn how to generate and view system reports. These reports are available in the **IBM Tivoli Storage Productivity Center > Reporting > System Reports** and **Data Manager for Databases > My Reports > System Reports** nodes in the navigation tree. Data for system reports are gathered every time that monitoring jobs are run against your storage resources.

### Generating and viewing system reports

Learn how to generate system reports, drill down through and filter their data, and generate charts.


To display a report, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center > My Reports > System Reports > <manager>**.
2. Click the name of the report you want to view. The report will appear in the content pane to the right. To regenerate a report, select the report name, right-click, and select **Regenerate Report**.


Many of the rows that are presented in the generated system reports enable you to:

- drill down into a row to get more detailed information
- view report information in a chart format

**To drill down for more detailed report information:**

- Click  on the desired row. The drill down report is displayed and a new tab representing that report is added to the tabs at the top of the report.

**To view report information in a chart format:**

- Click  on the desired row. The chart is displayed and a new tab representing that chart is added to the tabs at the top of the report.

You can customize the charts that you generate from system reports by using the Chart Customization window. You can access this window by right-clicking a system report chart and selecting **Customize this chart** from the pop-up menu. See “Customizing charts” on page 787 for information on how to customize a system report chart.

**To filter the data displayed on a report:**

Click the **Filter** button, which displays an Edit Filter window. In this window, specify criteria for records (for example, Capacity in Data Reports), operators (for

example, =) and values for each "column" (whatever you want to type). For example, in the Disk Space Summary report, you can specify only disks with a capacity greater than 70 GB.

If you change any element of a report definition (for example, the maximum number of rows in the report), you can save the new report definition (and regenerate it later) by clicking **File** → **Save** and then typing a name for the report. The report is saved with the name you specify under the *user\_id's Reports* node.

## System reports - Data


Use these predefined reports to view statistics about your storage resources.



### Access Time Summary

This report provides a summary of the number of files in your environment and when they were last accessed (for example, created or modified) during the last day, the last week, the last month, the last two months, the last three months, the last six months, the last nine months, the last year, and over a year.

The following table describes the columns on the Access Time Summary report. All the fields on the report are for display only.

|                                          |                                                                                  |
|------------------------------------------|----------------------------------------------------------------------------------|
| <b>Last Accessed &lt;= 1 day</b>         | The number and total size of the files accessed within the last day.             |
| <b>Last Accessed 1 day — 1 week</b>      | The number and total size of the files accessed between 1 day and 1 week ago.    |
| <b>Last Accessed 1 week — 1 month</b>    | The number and total size of the files accessed between 1 week and 1 month ago.  |
| <b>Last Accessed 1 month — 2 months</b>  | The number and total size the files accessed between 1 month and 2 months ago.   |
| <b>Last Accessed 2 months — 3 months</b> | The number and total size of the files accessed between 2 – 3 months ago.        |
| <b>Last Accessed 3 months — 6 months</b> | The number and total size of the files accessed between 3 – 6 months ago.        |
| <b>Last Accessed 6 months — 9 months</b> | The number and total size of the files accessed between 6 – 9 months ago.        |
| <b>Last Accessed 9 months — 1 year</b>   | The number and total size of the files accessed between 9 months and 1 year ago. |
| <b>Last Accessed &gt; 1 year</b>         | The number and total size of the files accessed over a year ago.                 |
| <b>Overall File Count</b>                | Total number of files across a network.                                          |
| <b>Overall File Size</b>                 | Total size of the space consumed by the files on a network.                      |
| <b>Average Age</b>                       | Average age of files on a network measured by days, hours, minutes, seconds.     |

- To view any of the following charts click  at the top of the report and select the corresponding option from the pop-up menu:
  - pie chart of storage space distribution for a selected row
  - pie chart of file counts for a selected row
  - pie chart of storage space distribution for all the rows on the report
  - pie chart of file counts for all the rows on a report
  - history chart of file counts
  - history chart of file count by percent
  - history chart of file sizes
  - history chart of files size by percent

- To view a pie chart of file access times, click  to left of a report row.
- To view detailed file access times for the computers on the network, click  to the left of a report row.

## Disk Space Summary




This report shows storage capacity for the whole network. You can drill down through this report to view storage capacity for each of the following levels: per disk, per computer, per cluster, per computer group, and per domain.

The following table describes the columns on the Disk Space Summary report. All the fields on the report are for display only.

---

|                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Disk Space</b>                             | Total storage capacity of the disks on the computers within a network                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Consumed File System Space</b>             | Amount of used storage space on the file systems within a network                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Available File System Space</b>            | amount of unused storage space on the file systems within a network                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Raw Volume Space</b>                       | Space on host-side logical volumes that is not occupied by file systems                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Unavailable Disk Space</b>                 | Amount of storage subsystem volume or virtual storage volume space that is dedicated to redundancy. This value is dependent on the storage subsystem SMI-S provider returning a valid RAID value that is used to determine the overhead. If the value is zero for a storage subsystem, the overhead cannot be calculated. This amount does not include storage space information from storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe. |
| <b>Available Disk Space</b>                   | Space assigned to a (monitored) host that is not part of any logical volume.                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Unknown Storage Subsystem Volume Space</b> | Volume space of unknown usage.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

---

- To view any of the following charts click  at the top of the report and select the corresponding option from the pop-up menu:
  - pie chart of network space distribution
  - pie chart of network row
  - history chart of the free space for a network
  - history chart of the percentage of free space for a network
  - history chart of the storage capacity on a network
- To view a pie chart of disk capacity, click  to the left of a report row.
- To view detailed disk capacity for the computers on a network, click  to the left of a report row.

## Access File Summary

This report provides overview information for files in your storage environment.

The Access File Summary report provides overview information for files by directory, directory group, file system, file system group, cluster, computer, computer group, domain, and for the entire network. Through this report you can



view historically the number of files for each resource in the report. The historical chart can be generated to show daily, weekly, or monthly history.


The following table describes the columns on the Access File Summary report. All the fields on the report are for display only.

---

|                          |                                                                    |
|--------------------------|--------------------------------------------------------------------|
| <b>File Size</b>         | Total size of the storage space consumed by the files on a network |
| <b>File Count</b>        | Total number of files on a network                                 |
| <b>Directory Count</b>   | Total number of directories on a network                           |
| <b>Average File Size</b> | Average storage space consumed by each of the files on a network   |
| <b>File System Space</b> | Total storage capacity of the files on a network                   |

---

- To view any of the following charts click  at the top of the report and select the corresponding option from the pop-up menu:
  - History chart of storage space usage for a selected row
  - History chart of file counts for a selected row
- To view a history chart of the storage space consumed by files on a network, click  to the left of a report row.

- To view detailed file storage usage for the computers on the network, click  to the left of a report row.

### Storage Access Times

This report shows when files were last accessed and how long ago they were accessed.

The following table describes the columns on the Storage Access Times report. All the fields on the report are for display only.


---

|                                          |                                                                                                                                                                      |
|------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Computer</b>                          | Name of a computer against which the report was run                                                                                                                  |
| <b>Last Accessed &lt;= 1 day</b>         | Number of files that were accessed within the last 24 hours and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)              |
| <b>Last Accessed 1 day - 1 week</b>      | Number of files that were accessed between 1 day to 1 week previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)      |
| <b>Last Accessed 1 week - 1 month</b>    | Number of files that were accessed between 1 week to 1 month previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)    |
| <b>Last Accessed 1 month - 2 months</b>  | Number of files that were accessed between 1 month to 2 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)  |
| <b>Last Accessed 2 months - 3 months</b> | Number of files that were accessed between 2 months to 3 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes) |
| <b>Last Accessed 3 months - 6 months</b> | Number of files that were accessed between 3 months to 6 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes) |

---

|                                          |                                                                                                                                                                      |
|------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Last Accessed 6 months - 9 months</b> | Number of files that were accessed between 6 months to 9 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes) |
| <b>Last Accessed 9 months - 1 year</b>   | Number of files that were accessed between 9 months to 1 year previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)   |
| <b>Last Accessed &gt; 1 year</b>         | Number of files that were accessed over one year previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)                |
| <b>Overall File Count</b>                | Total of all the counts                                                                                                                                              |
| <b>Overall File Size</b>                 | Total of all the sizes                                                                                                                                               |
| <b>Average Age</b>                       | Average time since the file was last accessed                                                                                                                        |

**To view graphical information about specific computers' access times:**

- Click  . A new tab displaying a chart for each selected computer is displayed.

**Disk Defects**

This report shows any disk defects on the computers being monitored.

The following table describes the columns on the Disk Defects report. All the fields on the report are for display only.


|                             |                                                                                    |
|-----------------------------|------------------------------------------------------------------------------------|
| <b>Device</b>               | Name of the computer containing a disk defect                                      |
| <b>Path</b>                 | Full path to the location of a disk defect                                         |
| <b>Primary Defects</b>      | Number of defects on the disk when it was new                                      |
| <b>Manufacturer</b>         | Name of the manufacturer for the disk drive containing the disk defect             |
| <b>Model</b>                | Model number of the disk drive containing a disk defect                            |
| <b>In Storage Subsystem</b> | Indicates Yes when a disk is an actual physical disk within the storage subsystem. |
| <b>RPM</b>                  | Speed that the disk turns (revolutions per minute)                                 |
| <b>Read Cache</b>           | Hardware caching turned on/off for reads                                           |
| <b>Write Cache</b>          | Hardware caching turned on/off for writes                                          |
| <b>Failure Predicted</b>    | Disk predicts a failure                                                            |
| <b>Probe Last Run</b>       | Date and time when the last probe was run                                          |
| <b>Discovered Time</b>      | Date and time when a disk defect was discovered                                    |
| <b>Disk Space</b>           | Storage capacity of the disk drive containing a defect                             |
| <b>Available Disk Space</b> | Amount of unused storage space on the disk drive containing a defect               |
| <b>Grown Defects</b>        | Number of defects detected by the disk since new                                   |
| <b>Recovered Errors</b>     | Number of errors (read, write, verify) from which the disk was able to recover     |

---

|                                                      |                                                                                                                             |
|------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| <b>Unrecovered Errors</b>                            | Number of errors (read, write, verify) from which the disk was not able to recover                                          |
| <b>Bytes Written</b>                                 | Number of bytes written to the disk                                                                                         |
| <b>Bytes Read</b>                                    | Number of bytes read from the disk                                                                                          |
| <b>Automatic Write Reallocation</b>                  | Disk attempts to move the data sector on the disk in the event of a write error                                             |
| <b>Desired Automatic Write Reallocation</b>          | Change automatic write allocation on the next probe                                                                         |
| <b>Automatic Read Reallocation</b>                   | Disk attempts to move the data sector on the disk in the event of a read error                                              |
| <b>Desired Automatic Read Reallocation</b>           | Change automatic read allocation on the next probe                                                                          |
| <b>Read Continuous (no error correction)</b>         | No error correction (setting is ON=no error correction)                                                                     |
| <b>Desired Read Continuous (no error correction)</b> | Change automatic read continuous on the next probe                                                                          |
| <b>Detectable</b>                                    | Indicates Yes when a storage volume is identified as being available to the storage subsystem following a successful probe. |
| <b>Is Virtual</b>                                    | Indicates Yes when a disk is identified as being available to the server or the storage subsystem.                          |
| <b>Mapped to Storage Subsystem</b>                   | Indicates Yes when a computer disk or a backend storage volume can be mapped to a storage subsystem.                        |

---

**To view more information for a specific computer:**

- Click  . Detailed information regarding the computer will be available in three tabs:
  - **General tab.** This tab displays general information about the computer, such as computer information, discovered date and time, disk allocation information, and so on.
  - **Latest Probe tab.** This tab displays information about the latest probe, and provides information such as sector size, probe last run, primary and grown defects, and so on.
  - **Probe History tab.** This tab displays probe history for the selected machine, and provides information such as probe times, record type, bytes written, and so on.

**Storage Modification Times**

This report shows information about files within the network that were modified with certain time frames.

The Storage Modification Times report provides information about files within the network that were modified:

- within the last 24 hours
- between 24 hours and one week previous
- between one week to one month previous
- between one month to two months previous
- between two months to three months previous
- between three months to six months previous

- between six months to nine months previous
- between nine months to one year previous
- more than one year previous


The following table describes the columns on the Storage Modification Times report. All the fields on the report are for display only.

---

|                                          |                                                                                                                                                                      |
|------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Last Modified &lt;= 1 day</b>         | Number of files that were modified in the last 24 hours and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)                  |
| <b>Last Modified 1 day - 1 week</b>      | Number of files that were modified between 1 day to 1 week previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)      |
| <b>Last Modified 1 week - 1 month</b>    | Number of files that were modified between 1 week to 1 month previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)    |
| <b>Last Modified 1 month - 2 months</b>  | Number of files that were modified between 1 month to 2 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)  |
| <b>Last Modified 2 months - 3 months</b> | Number of files that were modified between 2 months to 3 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes) |
| <b>Last Modified 3 months - 6 months</b> | Number of files that were modified between 3 months to 6 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes) |
| <b>Last Modified 6 months - 9 months</b> | Number of files that were modified between 6 months to 9 months previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes) |
| <b>Last Modified 9 months - 1 year</b>   | Number of files that were modified between 9 months to 1 year previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)   |
| <b>Last Modified &gt; 1 year</b>         | Number of files that were modified over a year previous and the total physical size of those files (measured in kilobytes, megabytes, or gigabytes)                  |
| <b>Overall File Count</b>                | Sum of the files                                                                                                                                                     |
| <b>Overall File Size</b>                 | Sum of the size of the files                                                                                                                                         |
| <b>Average Age</b>                       | Average age since files were modified                                                                                                                                |

---

**To view graphical information about storage modification times:**

- Click  . A new tab displaying a chart for the network is displayed.

### **Most At Risk Files**

This report provides information on files that have been modified the longest time ago and have not yet been backed-up or archived since they were modified.

The following table describes the columns on the Most At Risk Files report. All the fields on the report are for display only.

---

|                          |                                                      |
|--------------------------|------------------------------------------------------|
| <b>Modification Time</b> | Date and time when an at risk file was last modified |
| <b>Computer</b>          | Name of the computer where an at risk file is stored |
| <b>Filesystem</b>        | Filesystem where an at risk file is stored           |

---




---

|                      |                                                                                   |
|----------------------|-----------------------------------------------------------------------------------|
| <b>Path</b>          | Full path to the location of an at risk file                                      |
| <b>Filename</b>      | The name of the at risk file                                                      |
| <b>Owner</b>         | ID of the user who owns an at risk file                                           |
| <b>OS User Group</b> | OS User group of the user who owns an at risk file                                |
| <b>Logical Size</b>  | Logical size of an at risk file                                                   |
| <b>Physical Size</b> | Physical size of an at risk file (measured in kilobytes, megabytes, or gigabytes) |
| <b>Access Time</b>   | Date and time when an at risk file was last accessed                              |
| <b>Create Time</b>   | Date and time when an at risk file was created                                    |

---

**To view more information about a specific file:**

- Click  . Detailed information regarding the file will be displayed in a **File Detail** tab. The **File Detail** tab displays more information for the specific file and provides information such as file ownership, modification and creation time, and any file attributes that have been set.

**Most Obsolete Files**

This report provides information on files that have not been accessed or modified for the longest period of time.


The following table describes the columns on the Most Obsolete Files report. All the fields on the report are for display only.

---

|                          |                                                                                    |
|--------------------------|------------------------------------------------------------------------------------|
| <b>Access Time</b>       | Date and time when an obsolete file was last accessed                              |
| <b>Computer</b>          | Name of the computer where an obsolete file is stored                              |
| <b>Filesystem</b>        | Filesystem where an obsolete file is stored                                        |
| <b>Path</b>              | Full path to the location of an obsolete file                                      |
| <b>Filename</b>          | File name of an obsolete file                                                      |
| <b>Owner</b>             | ID of the user who owns an obsolete file                                           |
| <b>OS User Group</b>     | OS User Group to which the owner of the obsolete file belongs                      |
| <b>Logical Size</b>      | Logical size of an obsolete file                                                   |
| <b>Physical Size</b>     | Physical size of an obsolete file (measured in kilobytes, megabytes, or gigabytes) |
| <b>Modification Time</b> | Date and time when an obsolete file was last modified                              |
| <b>Create Time</b>       | Date and time when an obsolete file was created                                    |

---

**To view more information for a specific file:**

- Click  . Detailed information regarding the file will be displayed in a **File Detail** tab. The **File Detail** tab displays more information for the specific file and provides information such as file ownership, modification and creation time, and any file attributes that have been set.

## Oldest Orphaned Files

This report provides information on files that have the oldest creation date and no longer have the owners registered as users on the computer or network.


The following table describes the columns on the Oldest Orphaned Files report. All the fields on the report are for display only.

---

|                                       |                                                                                                                                                                           |
|---------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Access Time</b>                    | Date and time when an orphaned file was last accessed                                                                                                                     |
| <b>Computer</b>                       | Name of the computer where an orphaned file is stored                                                                                                                     |
| <b>Filesystem</b>                     | Filesystem where an orphaned file is stored                                                                                                                               |
| <b>Path<br/>Filename</b>              | Full path to the location of an orphaned file<br>File name of an orphaned file                                                                                            |
| <b>Owner</b>                          | Operating system internal ID of the user who owned the orphaned file. This is the internal ID the operating system uses to identify the user, and <i>not</i> the user ID. |
| <b>OS User Group<br/>Logical Size</b> | OS User group to which the owner of an orphaned file belongs<br>Logical size of an orphaned file                                                                          |
| <b>Physical Size</b>                  | Physical size of an orphaned file (measured in kilobytes, megabytes, or gigabytes)                                                                                        |
| <b>Modification Time</b>              | Date and time when an orphaned file was last modified                                                                                                                     |
| <b>Create Time</b>                    | Data and time when an orphaned file was created                                                                                                                           |

---

### To view more information for a specific file:

- Click  . Detailed information regarding the file will be displayed in a **File Detail** tab. The **File Detail** tab displays information about a specific file such as file ownership, modification and creation time, and any file attributes that have been set.

## User Quota Violations

This report shows which users have violated a Data Manager storage usage quota.

The following table describes the columns on the User Quota Violations report. All the fields on the report are for display only.

---

|                             |                                                      |
|-----------------------------|------------------------------------------------------|
| <b>Quota Creator</b>        | Creator of a quota                                   |
| <b>Quota Name</b>           | Name of a quota                                      |
| <b>Computer</b>             | Computer that violated a quota                       |
| <b>Filesystem</b>           | File system that violated a quota                    |
| <b>User Name</b>            | ID of the user that violated a quota                 |
| <b>Space Used</b>           | Total amount of space used                           |
| <b>Threshold</b>            | Threshold defined by a quota                         |
| <b>% of Threshold</b>       | Percentage of quota; always > 100                    |
| <b>Number of Violations</b> | Number of violations by a user                       |
| <b>Quota Type</b>           | Type of quota: Network Wide, Filesystem, or Computer |

---

## Computer Storage Availability

This report provides information about the availability of computers within your environment.


If the computer is not available (determined by a ping), the storage is likely not available for use. The following table describes the fields on the Storage Availability report. All the fields on the report are for display only.

---

|                          |                                                               |
|--------------------------|---------------------------------------------------------------|
| <b>Computer</b>          | Name of the computer                                          |
| <b>% Available</b>       | Percentage of successful pings                                |
| <b>Transmitted</b>       | Number of pings transmitted                                   |
| <b>Received</b>          | Number of responses from the ping                             |
| <b>Damaged</b>           | Number of responses with errors                               |
| <b>Avg Transit Time</b>  | Average time it took to send and receive response from a ping |
| <b>Max Transit Time</b>  | Maximum time it took to send and receive response from a ping |
| <b>Min Transit Time</b>  | Minimum time it took to send and receive response from a ping |
| <b>Report Start Time</b> | Start time of information in this row                         |
| <b>Report End Time</b>   | End time of information                                       |

---

**To view the ping history for a specific computer:**

- Click  . A new tab displaying ping history information for the selected computer is displayed.

**Computer Disk Space**

This report provides storage capacity information about each computer within your environment.

The following table describes the columns on the Computer Disk Space report. All the fields on the report are for display only.

---

|                             |                                                                                                  |
|-----------------------------|--------------------------------------------------------------------------------------------------|
| <b>Computer</b>             | Name of a computer against which the report was run                                              |
| <b>Disk Space</b>           | Total storage capacity for a computer                                                            |
| <b>Available Disk Space</b> | Amount of unused storage space on a computer (not in file systems seen by this operating system) |

---

|                                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Owned Disk Space</b>                                | <p>Total amount of unique disk storage space for computers in a cluster that is owned by those computers. If Tivoli Storage Productivity Center determines that a disk is configured for a different host, that disk's capacity is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:</p> <ul style="list-style-type: none"> <li>• Computer disks on computers that have not been probed.</li> <li>• Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.</li> </ul> <p><b>Note:</b></p> <ul style="list-style-type: none"> <li>• If the computer is an IBM SAN File System client, this value does not include the capacity of any volumes that are visible to it that are owned by the SAN File System.</li> <li>• If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and reads from the host ID. If this does not agree with the local host ID, the disk is classified as belonging to a different (Solaris) host.</li> </ul>  |
| <b>Owned Available Disk Space</b>                      | <p>Total amount of unique disk storage space for computers in a cluster that is not allocated to any logical volume within the computers and is owned by the computers. If Data Manager determines that a disk is configured for a different computer, that disk's unallocated space is not counted in this column. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:</p> <ul style="list-style-type: none"> <li>• Computer disks on computers that have not been probed.</li> <li>• Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from the computer.</li> </ul> <p>If the computer is an IBM SAN File System client, then this value excludes the unallocated space on any storage subsystem volumes visible to it that are owned by the SAN File System.</p> <p>If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product will examine the Veritas disk-label and read off the host-ID. If this does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.</p> |
| <b>Non-Fibre Channel Attached Disk Space</b>           | <p>Total amount of unique disk space on the computers in a cluster that is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:</p> <ul style="list-style-type: none"> <li>• Computer disks that reside behind a fibre channel port.</li> <li>• Computer disks on a computer that has not been probed.</li> <li>• Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Non-Fibre Channel Attached Available Disk Space</b> | <p>Total amount of disk storage space on the computers in a cluster that is not allocated to any logical volume within the computers and is not attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:</p> <ul style="list-style-type: none"> <li>• Computer disks that reside behind a fibre channel port.</li> <li>• Computer disks on a computer that has not been probed.</li> <li>• Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

|                                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Fibre Channel Attached Disk Space</b>           | Total amount of unique disk space on the computers in a cluster that is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column: <ul style="list-style-type: none"> <li>• Computer disks that do not reside behind a fibre channel port.</li> <li>• Computer disks that reside behind a fibre channel port that is not identified by a Storage Resource agent or Data agent.</li> <li>• Computer disks on a computer that has not been probed.</li> <li>• Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.</li> </ul>                                                                  |
| <b>Fibre Channel Attached Available Disk Space</b> | Total amount of disk storage space on the computers in a cluster that is not allocated to any logical volume within the computers and is attached through a fibre channel port. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column: <ul style="list-style-type: none"> <li>• Computer disks that do not reside behind a fibre channel port.</li> <li>• Computer disks that reside behind a fibre channel port that is not identified by a Storage Resource agent or Data agent.</li> <li>• Computer disks on a computer that has not been probed.</li> <li>• Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.</li> </ul> |
| <b>OS Type</b>                                     | Operating system running on a computer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Network Address</b>                             | Network address of a computer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>IP Address</b>                                  | IP address of a computer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Time Zone</b>                                   | Time zone in which a computer is running                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>CPU Architecture</b>                            | Instruction set architecture on the machine where an agent is running.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## Available File System Space

This report shows the total amount of unused storage across a network.

The following table describes the columns on the Total Freespace report. All the fields on the report are for display only.

|                                            |                                                                |
|--------------------------------------------|----------------------------------------------------------------|
| <b>Available File System Space</b>         | Total amount of available storage space available on a network |
| <b>Percent Available File System Space</b> | Percentage of total space that is unused on a network          |
| <b>Consumed File System Space</b>          | Amount of used storage space on a network                      |
| <b>File System Space</b>                   | Total amount (capacity) of storage space on a network          |
| <b>File Count</b>                          | Total number of files on a network                             |
| <b>Directory Count</b>                     | Total number of directories on a network                       |
| <b>Percent Free Inodes</b>                 | Percent of free Inodes on a network                            |
| <b>Used Inodes</b>                         | Number of used Inodes on a network                             |
| <b>Free Inodes</b>                         | Number of free Inodes on a network                             |

## User Space Usage

This report provides information about storage statistics related to users within your environment.

The following table describes the columns on the User Space Usage report. All the fields on the report are for display only.

---

|                            |                                                                       |
|----------------------------|-----------------------------------------------------------------------|
| <b>User Name</b>           | ID of a user                                                          |
| <b>File Size</b>           | Total amount of space used by a user                                  |
| <b>File Count</b>          | Number of files owned or created by a user                            |
| <b>Directory Count</b>     | Number of directories owned or created by a user                      |
| <b>Largest File</b>        | Largest file owned by a user                                          |
| <b>2nd Largest File</b>    | Second-largest file owned by a user                                   |
| <b>Files &lt; 1KB</b>      | Number and total space usage of files under 1KB in size               |
| <b>Files 1KB - 10KB</b>    | Number and total space usage of files between 1KB and 10KB in size    |
| <b>Files 10KB - 100KB</b>  | Number and total space usage of files between 10KB and 100KB in size  |
| <b>Files 100KB - 1MB</b>   | Number and total space usage of files between 100KB and 1MB in size   |
| <b>Files 1MB - 10MB</b>    | Number and total space usage of files between 1MB and 10MB in size    |
| <b>Files 10MB - 100MB</b>  | Number and total space usage of files between 10MB and 100MB in size  |
| <b>Files 100MB - 500MB</b> | Number and total space usage of files between 100MB and 500MB in size |
| <b>Files &gt; 500MB</b>    | Number and total space usage of files over 500MB in size              |

---

## Wasted Space

This report provides information about storage statistics on non-OS files not accessed in the first year and orphan files.


The following table describes the columns on the Wasted Space-File Summary by Computer report. All the fields on the report are for display only.

---

|                          |                                                             |
|--------------------------|-------------------------------------------------------------|
| <b>File Size</b>         | Total amount of space used by the obsolete and orphan files |
| <b>File Count</b>        | Total number of obsolete and orphan files                   |
| <b>Directory Count</b>   | Total number of orphan directories                          |
| <b>Average File Size</b> | Average size of obsolete and orphan files                   |
| <b>File System Space</b> | Amount of storage space available on a storage subsystem.   |

---

### To view more information for a specific file:

- Click  . Detailed information per computer will be displayed in a **Network** tab. The **Network** tab displays file summary by computer.

## System reports - Data Manager for Databases

Use these predefined reports to view statistics about your RDBMS storage resources. These reports are available in the **Data Manager for Databases > My Reports > System Reports** node of the navigation tree.

### All DBMS - Database Storage By Computer

This report shows storage information about the databases in your environment sorted by the computers on which they are stored.

---

|                        |                                                                       |
|------------------------|-----------------------------------------------------------------------|
| <b>Computer</b>        | Name of the computer where the databases are located                  |
| <b>Total Size</b>      | Amount of space consumed by the databases on the computers            |
| <b>File Capacity</b>   | Storage capacity of the data files within the databases               |
| <b>File Free space</b> | Amount of free space available on the data files within the databases |
| <b>DB-TS Count</b>     | Number tablespaces associated with the databases                      |
| <b>File Count</b>      | Number of data files associated with the tablespaces in the databases |
| <b>Log File Count</b>  | Number of log files associated with the databases                     |

---

### All DBMS - User Database Space Usage

This report shows information about the RDBMS storage space consumed by users across the network.

---

|                                   |                                                                                                    |
|-----------------------------------|----------------------------------------------------------------------------------------------------|
| <b>User Name</b>                  | Name of a user upon whose storage usage information was collected                                  |
| <b>Total Size</b>                 | Amount of space consumed by the objects owned by the user                                          |
| <b>Empty Used Space</b>           | Amount of empty used space residing on the objects owned by the user                               |
| <b>Table Count</b>                | Number of tables owned by the user                                                                 |
| <b>Index Count</b>                | Number of indexes owned by the user                                                                |
| <b>Other Count</b>                | Number of other objects owned by the user                                                          |
| <b>Largest Object</b>             | Size of the largest object owned by the user                                                       |
| <b>2nd Largest Object</b>         | Size of the second largest object owned by the user                                                |
| <b>Segment Size &lt; 50KB</b>     | Number and total space usage of objects owned by the user that are under 50KB in size              |
| <b>Segment Size 50KB - 500KB</b>  | Number and total space usage of objects owned by the user that are between 50KB and 500KB in size  |
| <b>Segment Size 500KB - 1Mb</b>   | Number and total space usage of objects owned by the user that are between 500KB and 1MB in size   |
| <b>Segment Size 1Mb - 10Mb</b>    | Number and total space usage of objects owned by the user that are between 1MB and 10MB in size    |
| <b>Segment Size 10Mb - 100Mb</b>  | Number and total space usage of objects owned by the user that are between 10MB and 100MB in size  |
| <b>Segment Size 100Mb - 500Mb</b> | Number and total space usage of objects owned by the user that are between 100MB and 500MB in size |
| <b>Segment Size 500Mb - 1GB</b>   | Number and total space usage of objects owned by the user that are between 500MB and 1GB in size   |
| <b>Segment Size &gt; 1Gb</b>      | Number and total space usage of objects owned by the user that are over 1GB in size                |

---

### All DBMS - User Database Quota Violations

This report provides information on the users who violated any storage usage thresholds defined in quotas.

---

|                      |                                                         |
|----------------------|---------------------------------------------------------|
| <b>Quota Creator</b> | ID of the user who created the quota                    |
| <b>Quota Name</b>    | Name of the quota                                       |
| <b>Computer</b>      | Name of the computer where the quota violation occurred |

---

|                             |                                                                                                                       |
|-----------------------------|-----------------------------------------------------------------------------------------------------------------------|
| <b>RDBMS type</b>           | RDBMS to which an instance belongs, for example, Oracle, SQL Server, Sybase, or UDB                                   |
| <b>Instance</b>             | Name/SID of the Oracle Instance, SQL Server server, Sybase server, or UDB Instance where the quota violation occurred |
| <b>Database-Tablespace</b>  | Name of the database where the tablespace resides and name of the tablespace whose quota was violated                 |
| <b>User Name</b>            | ID of the user who violated the quota                                                                                 |
| <b>Space Used</b>           | Amount of space consumed by the user                                                                                  |
| <b>Threshold</b>            | Threshold of used space defined within the quota                                                                      |
| <b>% of Threshold</b>       | Percentage of space by which the user violated the quota                                                              |
| <b>Number of Violations</b> | Number of violations for the user                                                                                     |
| <b>Quota Type</b>           | The type of quota: Database-Tablespace, Instance, or Network-wide                                                     |

### Oracle - Database Storage By Computer

This report shows storage information about the databases in your environment sorted by the computers on which they are stored.

|                             |                                                                       |
|-----------------------------|-----------------------------------------------------------------------|
| <b>Computer</b>             | Name of the computer where the databases are located                  |
| <b>Total Size</b>           | Amount of space consumed by the databases on the computers            |
| <b>Data File Capacity</b>   | Storage capacity of the data files within the databases               |
| <b>Data File Free space</b> | Amount of free space available on the data files within the databases |
| <b>Tablespace Count</b>     | Number tablespaces associated with the databases                      |
| <b>Data File Count</b>      | Number of data files associated with the tablespaces in the databases |
| <b>Log File Count</b>       | Number of log files associated with the databases                     |

### Oracle - Total Database Freespace

This report provides information on the total free space for data files at a network-wide level, as well as the percentage of free space, the total used space, number of free extents, and number of data files on a network. This report is available for Oracle and Microsoft SQL Server.

|                                 |                                                                   |
|---------------------------------|-------------------------------------------------------------------|
| <b>Free Space</b>               | Amount of free space on all data files in the network.            |
| <b>Percent Free</b>             | Percent of free space available on the data files in the network. |
| <b>Used Space</b>               | Amount of used space on the data files in the network.            |
| <b>Total Size</b>               | Total size of data files in the network.                          |
| <b>Free Extents</b>             | Number of free extents on the data files in the network.          |
| <b>Coalesced Extents</b>        | Number of coalesced extents on the databases in the network.      |
| <b>Minimum Free Extent Size</b> | Minimum size of a free extent in a database.                      |
| <b>Maximum Free Extent Size</b> | Maximum size of a free extent in a database.                      |
| <b>Number Data Files</b>        | Number of data files on the databases in the network              |

### Oracle - User Database Space Usage

This report shows information about the RDBMS storage space consumed by users across the network.

|                            |                                                                      |
|----------------------------|----------------------------------------------------------------------|
| <b>User Name</b>           | Name of a user upon whose storage usage information was collected    |
| <b>Total Size</b>          | Amount of space consumed by the objects owned by the user            |
| <b>Empty Used Space</b>    | Amount of empty used space residing on the objects owned by the user |
| <b>Table Count</b>         | Number of tables owned by the user                                   |
| <b>Index Count</b>         | Number of indexes owned by the user                                  |
| <b>Other Count</b>         | Number of other objects owned by the user                            |
| <b>Largest Segment</b>     | Size of the largest segment owned by the user                        |
| <b>2nd Largest Segment</b> | Size of the second largest segment owned by the user                 |



|                                   |                                                                                                    |
|-----------------------------------|----------------------------------------------------------------------------------------------------|
| <b>Segment Size &lt; 50KB</b>     | Number and total space usage of objects owned by the user that are under 50KB in size              |
| <b>Segment Size 50KB - 500KB</b>  | Number and total space usage of objects owned by the user that are between 50KB and 500KB in size  |
| <b>Segment Size 500KB - 1Mb</b>   | Number and total space usage of objects owned by the user that are between 500KB and 1MB in size   |
| <b>Segment Size 1Mb - 10Mb</b>    | Number and total space usage of objects owned by the user that are between 1MB and 10MB in size    |
| <b>Segment Size 10Mb - 100Mb</b>  | Number and total space usage of objects owned by the user that are between 10MB and 100MB in size  |
| <b>Segment Size 100Mb - 500Mb</b> | Number and total space usage of objects owned by the user that are between 100MB and 500MB in size |
| <b>Segment Size 500Mb - 1GB</b>   | Number and total space usage of objects owned by the user that are between 500MB and 1GB in size   |
| <b>Segment Size &gt; 1Gb</b>      | Number and total space usage of objects owned by the user that are over 1GB in size                |

### Oracle - Segments With Wasted Space

This report provides information on the segments containing allocated space that is currently empty/not being used. Use this information to gather information that can help you discover space that can be reclaimed and allocated to other objects.

|                          |                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Empty Used Space</b>  | Amount of empty used space within a segment (table, index, etc.)                                                                                                                                                                                                                                                                                                       |
| <b>Segment Creator</b>   | Owner of a segment                                                                                                                                                                                                                                                                                                                                                     |
| <b>Segment Name</b>      | Name of a segment                                                                                                                                                                                                                                                                                                                                                      |
| <b>Computer</b>          | Name of the computer on which the segment's Instance resides                                                                                                                                                                                                                                                                                                           |
| <b>Instance</b>          | SID of an Oracle Instance                                                                                                                                                                                                                                                                                                                                              |
| <b>Database</b>          | Name of the database to which the segment belongs                                                                                                                                                                                                                                                                                                                      |
| <b>Tablespace</b>        | Name of the tablespace to which the segment belongs                                                                                                                                                                                                                                                                                                                    |
| <b>Partition</b>         | Partition on which the segment is stored                                                                                                                                                                                                                                                                                                                               |
| <b>Segment Type</b>      | Type of the segment, including: <ul style="list-style-type: none"> <li>• TABLE</li> <li>• TABLE PARTITION</li> <li>• TABLE SUBPARTITION</li> <li>• NESTED TABLE</li> <li>• CLUSTER</li> <li>• INDEX</li> <li>• INDEX PARTITION</li> <li>• INDEX SUBPARTITION</li> <li>• LOBINDEX</li> <li>• LOBSEGMENT</li> <li>• LOB PARTITION</li> <li>• LOB SUBPARTITION</li> </ul> |
| <b>Parent Type</b>       | Subset of the segment type                                                                                                                                                                                                                                                                                                                                             |
| <b>Parent Creator</b>    | Owner of a segment                                                                                                                                                                                                                                                                                                                                                     |
| <b>Parent Name</b>       | Name of a segment                                                                                                                                                                                                                                                                                                                                                      |
| <b>Total Size</b>        | Amount of space allocated to a segment                                                                                                                                                                                                                                                                                                                                 |
| <b>Number of Extents</b> | Number of extents allocated to a segment                                                                                                                                                                                                                                                                                                                               |
| <b>Freelist Blocks</b>   | Number of blocks on the freelist chain                                                                                                                                                                                                                                                                                                                                 |
| <b>Initial Extent</b>    | Size of the first extent allocated to a segment                                                                                                                                                                                                                                                                                                                        |
| <b>Next Extent</b>       | Amount of space Oracle will retrieve when allocating another extent                                                                                                                                                                                                                                                                                                    |
| <b>Maximum Extents</b>   | Maximum number of extents that Oracle would allocate to an object                                                                                                                                                                                                                                                                                                      |
| <b>Percent Increase</b>  | Percent increase in size Oracle will allocate for the next extent                                                                                                                                                                                                                                                                                                      |

## SqlServer - Database Storage By Computer

This report shows storage information about the databases in your environment sorted by the computers on which they are stored.

---

|                             |                                                                                                           |
|-----------------------------|-----------------------------------------------------------------------------------------------------------|
| <b>Computer</b>             | Name of the computer on which SQL Server is running                                                       |
| <b>Total Size</b>           | Total amount of storage space allocated to SQL Server on a computer                                       |
| <b>Data File Capacity</b>   | Amount of allocated storage space for data files that is unused on a computer where SQL Server is running |
| <b>Data File Free space</b> | Amount of free space available on the data files associated with SQL Server on a computer                 |
| <b>Log File Capacity</b>    | Amount of allocated storage space for log files that is unused on a computer where SQL Server is running  |
| <b>Log File Free space</b>  | Amount of free space available on the log files associated with SQL Server on a computer                  |
| <b>Database Count</b>       | Number of databases associated with SQL Server on a computer                                              |
| <b>Data File Count</b>      | Number of data files associated with SQL Server on a computer                                             |
| <b>Log File Count</b>       | Number of log files associated with SQL Server on a computer                                              |

---

## SqlServer - Total Database Freespace

This report provides information on the total free space for data files at a network-wide level, as well as the percentage of free space, the total used space, number of free extents, and number of data files on a network. This report is available for Oracle and Microsoft SQL Server.

---

|                          |                                                                   |
|--------------------------|-------------------------------------------------------------------|
| <b>Data Free Space</b>   | Amount of free space on all data files in the network.            |
| <b>Percent Free</b>      | Percent of free space available on the data files in the network. |
| <b>Data Used Space</b>   | Amount of used space on the data files in the network.            |
| <b>Data Size</b>         | Total size of data files in the network.                          |
| <b>Log Free Space</b>    | Amount of free space.                                             |
| <b>Log Percent Used</b>  | Percent of used space.                                            |
| <b>Log Used Space</b>    | Amount of space used by the log.                                  |
| <b>Log Size</b>          | Size of the log.                                                  |
| <b>Number Data Files</b> | Number of data files on the databases in the network              |

---

## SqlServer - User Database Space Usage

This report shows information about the RDBMS storage space consumed by users across the network.

---

|                                 |                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------|
| <b>User Name</b>                | Name of a user upon whose storage usage information was collected                                 |
| <b>Total Size</b>               | Amount of space consumed by the objects owned by the user                                         |
| <b>Empty Used Space</b>         | Amount of empty used space residing on the objects owned by the user                              |
| <b>Table Count</b>              | Number of tables owned by the user                                                                |
| <b>Index Count</b>              | Number of indexes owned by the user                                                               |
| <b>Other Count</b>              | Number of other objects owned by the user                                                         |
| <b>Largest Object</b>           | Size of the largest object owned by the user                                                      |
| <b>2nd Largest Object</b>       | Size of the second largest object owned by the user                                               |
| <b>Object Size &lt; 50KB</b>    | Number and total space usage of objects owned by the user that are under 50KB in size             |
| <b>Object Size 50KB - 500KB</b> | Number and total space usage of objects owned by the user that are between 50KB and 500KB in size |
| <b>Object Size 500KB - 1Mb</b>  | Number and total space usage of objects owned by the user that are between 500KB and 1MB in size  |
| <b>Object Size 1Mb - 10Mb</b>   | Number and total space usage of objects owned by the user that are between 1MB and 10MB in size   |

---

|                                  |                                                                                                    |
|----------------------------------|----------------------------------------------------------------------------------------------------|
| <b>Object Size 10Mb - 100Mb</b>  | Number and total space usage of objects owned by the user that are between 10MB and 100MB in size  |
| <b>Object Size 100Mb - 500Mb</b> | Number and total space usage of objects owned by the user that are between 100MB and 500MB in size |
| <b>Object Size 500Mb - 1GB</b>   | Number and total space usage of objects owned by the user that are between 500MB and 1GB in size   |
| <b>Object Size &gt; 1Gb</b>      | Number and total space usage of objects owned by the user that are over 1GB in size                |

## Sybase - Database Storage By Computer

This report shows storage information about the databases in your environment sorted by the computers on which they are stored.

|                                 |                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------|
| <b>Computer</b>                 | Name of the computer on which Sybase is running                                                          |
| <b>Total Device Size</b>        | Total amount of storage space allocated to Sybase on a computer                                          |
| <b>Mirrored Device Size</b>     | Total storage capacity of the mirrored device for an instance                                            |
| <b>Device Free space</b>        | Amount of free space available on the data files associated with Sybase on a computer                    |
| <b>Data Fragment Capacity</b>   | Amount of allocated storage space for log files that is unused on a computer where Sybase is running     |
| <b>Data Fragment Free space</b> | Amount of free space available on the data fragments associated with Sybase on a computer                |
| <b>Log Fragment Capacity</b>    | Amount of allocated storage space for log fragments that is unused on a computer where Sybase is running |
| <b>Log Fragment Free space</b>  | Amount of free space available on the log files associated with Sybase on a computer                     |
| <b>Device Count</b>             | Number of devices associated with Sybase on a computer                                                   |
| <b>Database Count</b>           | Number of databases associated with Sybase on a computer                                                 |
| <b>Fragment Count</b>           | Number of fragments associated with Sybase on a computer                                                 |
| <b>Log File Count</b>           | Number of log files associated with Sybase on a computer                                                 |

## Sybase - Total Device Freespace

This report provides information on the total free space for Sybase devices at a network-wide level, as well as the percentage of free space, the total used space, number of fragments, and number of devices.

|                         |                                                                |
|-------------------------|----------------------------------------------------------------|
| <b>Free Space</b>       | Amount of free space on all the devices in the network         |
| <b>Percent Free</b>     | Percent of free space available on the devices in the network  |
| <b>Used Space</b>       | Amount of used space on the devices in the network             |
| <b>Total Size</b>       | Total size of devices in the network                           |
| <b>Number Fragments</b> | Number of fragments associated with the devices in the network |
| <b>Number Devices</b>   | Number of devices in the network                               |

## Sybase - Total Database Freespace

This report provides information on the total free space for data files at a network-wide level, as well as the percentage of free space, the total used space, number of free extents, and number of data files on a network. This report is available for Oracle and Microsoft SQL Server.

|                        |                                                                   |
|------------------------|-------------------------------------------------------------------|
| <b>Data Free Space</b> | Amount of free space on all data files in the network.            |
| <b>Percent Free</b>    | Percent of free space available on the data files in the network. |
| <b>Data Used Space</b> | Amount of used space on the data files in the network.            |
| <b>Data Size</b>       | Total size of data files in the network.                          |
| <b>Log Free Space</b>  | Amount of free space available to a database log                  |

|                         |                                                                |
|-------------------------|----------------------------------------------------------------|
| <b>Log Percent Used</b> | Percentage of allocated space that a database log is consuming |
| <b>Log Used Space</b>   | Amount of storage space consumed by a database log             |
| <b>Log Size</b>         | Total amount of space allocated to a database log              |
| <b>Number Fragments</b> | nNumber of fragments associated with a database                |

## Sybase - User Database Space Usage

This report shows information about the RDBMS storage space consumed by users across the network.

|                                  |                                                                                                    |
|----------------------------------|----------------------------------------------------------------------------------------------------|
| <b>User Name</b>                 | Name of a user upon whose storage usage information was collected                                  |
| <b>Total Size</b>                | Amount of space consumed by the objects owned by the user                                          |
| <b>Empty Used Space</b>          | Amount of empty used space residing on the objects owned by the user                               |
| <b>Table Count</b>               | Number of tables owned by the user                                                                 |
| <b>Index Count</b>               | Number of indexes owned by the user                                                                |
| <b>Other Count</b>               | Number of other objects owned by the user                                                          |
| <b>Largest Object</b>            | Size of the largest object owned by the user                                                       |
| <b>2nd Largest Object</b>        | Size of the second largest object owned by the user                                                |
| <b>Object Size &lt; 50KB</b>     | Number and total space usage of objects owned by the user that are under 50KB in size              |
| <b>Object Size 50KB - 500KB</b>  | Number and total space usage of objects owned by the user that are between 50KB and 500KB in size  |
| <b>Object Size 500KB - 1Mb</b>   | Number and total space usage of objects owned by the user that are between 500KB and 1MB in size   |
| <b>Object Size 1Mb - 10Mb</b>    | Number and total space usage of objects owned by the user that are between 1MB and 10MB in size    |
| <b>Object Size 10Mb - 100Mb</b>  | Number and total space usage of objects owned by the user that are between 10MB and 100MB in size  |
| <b>Object Size 100Mb - 500Mb</b> | Number and total space usage of objects owned by the user that are between 100MB and 500MB in size |
| <b>Object Size 500Mb - 1GB</b>   | Number and total space usage of objects owned by the user that are between 500MB and 1GB in size   |
| <b>Object Size &gt; 1Gb</b>      | Number and total space usage of objects owned by the user that are over 1GB in size                |

## UDB - Instance Storage By Computer

This report shows information about the UDB instances in your environment sorted by the computers on which they are stored.

|                             |                                                                               |
|-----------------------------|-------------------------------------------------------------------------------|
| <b>Computer</b>             | Name of the computer on which a UDB Instance(s) resides                       |
| <b>Total Size</b>           | Total amount of storage space allocated to UDB on a computer                  |
| <b>Container Capacity</b>   | Amount of storage space allocated to the containers that reside on a computer |
| <b>Container Free space</b> | Amount of free space available on the containers that reside on a computer    |
| <b>Log File Capacity</b>    | Amount of storage space allocated to the log files that reside on a computer  |
| <b>Tablespace Count</b>     | Number of tablespaces that reside on a computer                               |
| <b>Container Count</b>      | Number of containers that reside on a computer                                |
| <b>Log File Count</b>       | Number of log files that reside on a computer                                 |

## UDB - Total DMS Container Freespace

This report shows the total free space for the containers associated with DMS tablespaces on UDB instances within your environment.

|                     |                                                                           |
|---------------------|---------------------------------------------------------------------------|
| <b>Free Space</b>   | Amount of free space available on the DMS containers within a network     |
| <b>Percent Free</b> | Percentage of free space available on the DMS containers within a network |

|                          |                                                                         |
|--------------------------|-------------------------------------------------------------------------|
| <b>Used Space</b>        | Amount of storage space consumed on the DMS containers within a network |
| <b>Total Size</b>        | Total amount of space on the DMS containers within a network            |
| <b>Number Containers</b> | Number of DMS containers within a network                               |

## UDB - User Database Space Usage

This report shows information about the RDBMS storage space consumed by users across the network.

|                                  |                                                                                                    |
|----------------------------------|----------------------------------------------------------------------------------------------------|
| <b>User Name</b>                 | Name of a user upon whose storage usage information was collected                                  |
| <b>Total Size</b>                | Amount of space consumed by the objects owned by the user                                          |
| <b>Empty Used Space</b>          | Amount of empty used space residing on the objects owned by the user                               |
| <b>Table Count</b>               | Number of tables owned by the user                                                                 |
| <b>Index Count</b>               | Number of indexes owned by the user                                                                |
| <b>Other Count</b>               | Number of other objects owned by the user                                                          |
| <b>Largest Object</b>            | Size of the largest segment owned by the user                                                      |
| <b>2nd Largest Object</b>        | Size of the second largest segment owned by the user                                               |
| <b>Object Size &lt; 50KB</b>     | Number and total space usage of objects owned by the user that are under 50KB in size              |
| <b>Object Size 50KB - 500KB</b>  | Number and total space usage of objects owned by the user that are between 50KB and 500KB in size  |
| <b>Object Size 500KB - 1Mb</b>   | Number and total space usage of objects owned by the user that are between 500KB and 1MB in size   |
| <b>Object Size 1Mb - 10Mb</b>    | Number and total space usage of objects owned by the user that are between 1MB and 10MB in size    |
| <b>Object Size 10Mb - 100Mb</b>  | Number and total space usage of objects owned by the user that are between 10MB and 100MB in size  |
| <b>Object Size 100Mb - 500Mb</b> | Number and total space usage of objects owned by the user that are between 100MB and 500MB in size |
| <b>Object Size 500Mb - 1GB</b>   | Number and total space usage of objects owned by the user that are between 500MB and 1GB in size   |
| <b>Object Size &gt; 1Gb</b>      | Number and total space usage of objects owned by the user that are over 1GB in size                |

## System Reports - Disk Manager

Use these predefined reports to view performance information about storage subsystems.

The following system reports are available for Disk Manager .

### Array Performance

This report shows performance data for arrays. This report supports IBM System Storage DS8000, DS6000, and IBM TotalStorage Enterprise Storage Server systems only.

| Column    | Description                                                                                          |
|-----------|------------------------------------------------------------------------------------------------------|
| Subsystem | ID of the storage system                                                                             |
| Array     | The ID of the specified array.                                                                       |
| Time      | Date and time that the data was collected.                                                           |
| Interval  | The actual duration of the sample interval (in seconds) for the performance data that was collected. |

| Column                        | Description                                                                                                                                                                                           |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Disk Utilization Percentage   | The approximate utilization percentage of a rank over a specified time interval (the average percent of time that the disks associated with the array were busy).                                     |
| Sequential I/O Percentage     | Percentage of all I/O operations performed for an array over a specified time interval that were sequential operations.                                                                               |
| Backend Read I/O Rate         | Average number of I/O operations per second for read operations.                                                                                                                                      |
| Backend Write I/O Rate        | Average number of I/O operations per second for write operations.                                                                                                                                     |
| Total Backend I/O Rate        | Average number of I/O operations per second for read and write operations.                                                                                                                            |
| Backend Read Data Rate        | Average number of megabytes (2 <sup>20</sup> bytes) that were transferred for read operations.                                                                                                        |
| Backend Write Data Rate       | Average number of megabytes (2 <sup>20</sup> bytes) that were transferred for write operations.                                                                                                       |
| Total Backend Data Rate       | Average number of megabytes (2 <sup>20</sup> bytes) that were transferred for read and write operations.                                                                                              |
| Backend Read Response Time    | Average number of milliseconds that it took to respond to each read operation. For SAN Volume Controller models, this is the external response time of the managed disks (MDisks).                    |
| Backend Write Response Time   | Average number of milliseconds that it took to respond to each write operation. For SAN Volume Controller models, this is the external response time of the managed disks (MDisks).                   |
| Overall Backend Response Time | Average number of milliseconds that it took to respond to each I/O operation (read and write). For SAN Volume Controller models, this is the external response time of the managed disks (MDisks).    |
| Write-cache Delay Percentage  | Percentage of I/O operations that were delayed due to write-cache space constraints or other conditions. Write-cache Delay Percentage is the ratio of delayed I/O operations to total I/O operations. |
| Write-cache Delay I/O Rate    | Average number of I/O operations per second that were delayed because of write-cache space constraints or other conditions, for a particular component over an interval.                              |

### Controller Cache Performance

This report shows performance data for storage system controller caches. This report supports IBM System Storage DS8000, DS6000, DS5000, DS4000, and IBM TotalStorage Enterprise Storage Server systems only.

| Column                                | Description                                                                                                                                                  |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Subsystem                             | ID of the storage system.                                                                                                                                    |
| Controller                            | The ID of the controller.                                                                                                                                    |
| Time                                  | Date and time that the data was collected.                                                                                                                   |
| Interval                              | The actual duration of the sample interval (in seconds) for the performance data that was collected.                                                         |
| Read I/O Rate (normal)                | Average number of I/O operations per second for non-sequential read operations for a component over a specified time interval.                               |
| Read I/O Rate (sequential)            | Average number of I/O operations per second for sequential read operations for a component over a specified time interval.                                   |
| Read I/O Rate (overall)               | Average number of I/O operations per second for both sequential and non-sequential read operations for a component over a specified time interval.           |
| Write I/O Rate (normal)               | Average number of I/O operations per second for non-sequential write operations for a component over a specified time interval.                              |
| Write I/O Rate (sequential)           | Average number of I/O operations per second for sequential write operations for a component over a specified time interval.                                  |
| Write I/O Rate (overall)              | Average number of I/O operations per second for both sequential and non-sequential write operations for a component over a specified time interval.          |
| Total I/O Rate (normal)               | Average number of I/O operations per second for non-sequential read and write operations for a component over a specified time interval.                     |
| Total I/O Rate (sequential)           | Average number of I/O operations per second for sequential read and write operations for a component over a specified time interval.                         |
| Total I/O Rate (overall)              | Average number of I/O operations per second for both sequential and non-sequential read and write operations for a component over a specified time interval. |
| Cache Holding Time                    | Average cache holding time, in seconds, for I/O data in this subsystem controller (cluster). Shorter time periods indicate adverse performance.              |
| Read Cache Hits Percentage (overall)  | Percentage of cache hits for both sequential and non-sequential read operations for a component over a specified time interval.                              |
| Write Cache Hits Percentage (overall) | Percentage of cache hits for both sequential and non-sequential write operations for a component over a specified time interval.                             |

| Column                                    | Description                                                                                                                                                                                           |
|-------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Total Cache Hits Percentage (overall)     | Percentage of cache hits for both sequential and non-sequential read and write operations for a component over a specified time interval.                                                             |
| Read Cache Hits Percentage (normal)       | Percentage of cache hits for non-sequential read operations for a component over a specified time interval.                                                                                           |
| Write Cache Hits Percentage (normal)      | Percentage of cache hits for non-sequential write operations for a component over a specified time interval.                                                                                          |
| Total Cache Hits Percentage (normal)      | Percentage of cache hits for non-sequential read and write operations for a component over a specified time interval.                                                                                 |
| Read Cache Hits Performance (sequential)  | The sequential read cache allocated for the controller.                                                                                                                                               |
| Write Cache Hits Performance (sequential) | The sequential write cache allocated for the controller.                                                                                                                                              |
| Total Cache Hits Performance (sequential) | The total sequential cache percentage (read and write) allocated for the controller.                                                                                                                  |
| Record Mode Read I/O Rate                 | Average number of I/O operations per second for record mode read operations for a component over a specified time interval.                                                                           |
| Record Mode Read Cache Hits Percentage    | Percentage of cache hits for record mode read operations for a component over a specified time interval.                                                                                              |
| Disk to Cache I/O Rate                    | Average number of I/O operations (track transfers) per second for disk to cache transfers for a component over a specified time interval.                                                             |
| Cache to Disk I/O Rate                    | Average number of I/O operations (track transfers) per second for cache to disk transfers for a component over a specified time interval.                                                             |
| Write-cache Delay Percentage              | Percentage of I/O operations that were delayed due to write-cache space constraints or other conditions. Write-cache Delay Percentage is the ratio of delayed I/O operations to total I/O operations. |
| Write-cache Delay I/O Rate                | Average number of I/O operations per second that were delayed due to write-cache space constraints or other conditions for a component over a specified time interval.                                |

## Controller Performance

This report shows performance data for storage system controllers. This report supports IBM System Storage DS8000, DS6000, DS5000, DS4000, and IBM TotalStorage Enterprise Storage Server systems only.

| Column     | Description               |
|------------|---------------------------|
| Subsystem  | ID of the storage system. |
| Controller | The ID of the controller. |



| Column                   | Description                                                                                                                                                        |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Time                     | Date and time that the data was collected.                                                                                                                         |
| Interval                 | The actual duration of the sample interval (in seconds) for the performance data that was collected.                                                               |
| Read I/O Rate (overall)  | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for read operations for a component over a specified time interval.           |
| Write I/O Rate (overall) | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for write operations for a component over a specified time interval.          |
| Total I/O Rate (overall) | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for read and write operations for a component over a specified time interval. |
| Read Data Rate           | Average number of I/O operations per second for both sequential and non-sequential read operations for a component over a specified time interval.                 |
| Write Data Rate          | Average number of I/O operations per second for both sequential and non-sequential write operations for a component over a specified time interval.                |
| Total Data Rate          | Average number of I/O operations per second for both sequential and nonsequential read and write operations.                                                       |
| Read Response Time       | Average number of milliseconds that it took to service each read operation for a component over a specified time interval.                                         |
| Write Response Time      | Average number of milliseconds that it took to service each write operation for a component over a specified time interval.                                        |
| Overall Response Time    | Average number of milliseconds that it took to service each I/O operation (read and write) for a component over a specified time interval.                         |

## I/O Group Performance

This report shows performance data for I/O groups. This report supports IBM System Storage SAN Volume Controller systems only.

| Column                     | Description                                                                                          |
|----------------------------|------------------------------------------------------------------------------------------------------|
| Subsystem                  | ID of the storage system                                                                             |
| I/O Group                  | The ID of the specified input/output (I/O) group.                                                    |
| Time                       | Date and time that the data was collected.                                                           |
| Interval                   | The actual duration of the sample interval (in seconds) for the performance data that was collected. |
| CPU Utilization Percentage | Average utilization percentage of the CPUs.                                                          |

| <b>Column</b>                 | <b>Description</b>                                                                                                                                                 |
|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Read I/O Rate (overall)       | Average number of I/O operations per second for both sequential and non-sequential read operations for a component over a specified time interval.                 |
| Write I/O Rate (overall)      | Average number of I/O operations per second for both sequential and non-sequential write operations for a component over a specified time interval.                |
| Total I/O Rate (overall)      | Average number of I/O operations per second for both sequential and nonsequential read and write operations.                                                       |
| Read Data Rate                | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for read operations for a component over a specified time interval.           |
| Write Data Rate               | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for write operations for a component over a specified time interval.          |
| Total Data Rate               | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for read and write operations for a component over a specified time interval. |
| Read Response Time            | Average number of milliseconds that it took to service each read operation for a component over a specified time interval.                                         |
| Write Response Time           | Average number of milliseconds that it took to service each write operation for a component over a specified time interval.                                        |
| Overall Response Time         | Average number of milliseconds that it took to service each I/O operation (read and write) for a component over a specified time interval.                         |
| Port to Host Send I/O Rate    | Average number of exchanges (I/Os) per second sent to host computers by a component over a specified time interval.                                                |
| Port to Host Receive I/O Rate | Average number of exchanges (I/Os) per second received from host computers by a component over a specified time interval.                                          |
| Total Port to Host I/O Rate   | Average number of exchanges (I/Os) per second transmitted between host computers and a component over a specified time interval.                                   |
| Port to Disk Send I/O Rate    | Average number of exchanges (I/Os) per second sent to storage subsystems by a component over a specified time interval.                                            |
| Port to Disk Receive I/O Rate | Average number of exchanges (I/Os) per second received from storage subsystems by a component over a specified time interval.                                      |

| Column                         | Description                                                                                                                                           |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| Total Port to Disk I/O Rate    | Average number of exchanges (I/Os) per second transmitted between storage subsystems and a component over a specified time interval.                  |
| Port to Host Send Data Rate    | Average number of megabytes (2 <sup>20</sup> bytes) per second sent to host computers by a component over a specified time interval.                  |
| Port to Host Receive Data Rate | Average number of megabytes (2 <sup>20</sup> bytes) per second received from host computers by a component over a specified time interval.            |
| Total Port to Host Data Rate   | Average number of megabytes (2 <sup>20</sup> bytes) per second transmitted between host computers and a component over a specified time interval.     |
| Port to Disk Send Data Rate    | Average number of megabytes (2 <sup>20</sup> bytes) per second sent to storage subsystems by a component over a specified time interval.              |
| Port to Disk Receive Data Rate | Average number of megabytes (2 <sup>20</sup> bytes) per second received from storage subsystems by a component over a specified time interval.        |
| Total Port to Disk Data Rate   | Average number of megabytes (2 <sup>20</sup> bytes) per second transmitted between storage subsystems and a component over a specified time interval. |

## Managed Disk Group Performance

This report shows performance data for managed-disk (MDisk) groups. This report supports IBM System Storage SAN Volume Controller systems only.

| Column                 | Description                                                                                                   |
|------------------------|---------------------------------------------------------------------------------------------------------------|
| Subsystem              | ID of the storage system                                                                                      |
| Managed Disk Group     | The ID of the specified managed disk.                                                                         |
| Time                   | Date and time that the data was collected.                                                                    |
| Interval               | The actual duration of the sample interval (in seconds) for the performance data that was collected.          |
| Read Transfer Size     | Average number of KB per input/output (I/O) for read operations, for a particular component over an interval. |
| Write Transfer Size    | Average number of KB per I/O for write operations, for a particular component over an interval.               |
| Overall Transfer Size  | Average number of KB per I/O for read and write operations, for a particular component over an interval.      |
| Backend Read I/O Rate  | Average number of I/O operations per second for read operations.                                              |
| Backend Write I/O Rate | Average number of I/O operations per second for write operations.                                             |

| Column                  | Description                                                                 |
|-------------------------|-----------------------------------------------------------------------------|
| Total Backend I/O Rate  | Average number of I/O operations per second for read and write operations.  |
| Backend Read Data Rate  | Average number of Mbps that were transferred for read operations.           |
| Backend Write Data Rate | Average number of Mbps that were transferred for write operations.          |
| Total Backend Data Rate | Average number of Mbps that were transferred for read and write operations. |

## Node Cache Performance report

This report lists the most recently gathered performance data for IBM System Storage SAN Volume Controller nodes, which is aggregated to the node level from the per-node VDisk sample data.

| Column                                | Description                                                                                                                                         |
|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| Subsystem                             | ID of the storage system                                                                                                                            |
| Node                                  | The ID of the specified node.                                                                                                                       |
| Time                                  | Date and time that the data was collected.                                                                                                          |
| Interval                              | The actual duration of the sample interval (in seconds) for the performance data that was collected.                                                |
| Read Cache Hits Percentage (overall)  | Percentage of cache hits for both sequential and nonsequential read operations, for a particular component over an interval.                        |
| Write Cache Hits Percentage (overall) | Percentage of cache hits for both sequential and nonsequential write operations, for a particular component over an interval.                       |
| Total Cache Hits Percentage (overall) | Percentage of cache hits for both sequential and nonsequential read and write operations, for a particular component over an interval.              |
| Read I/O Rate (overall)               | Average number of I/O operations per second for both sequential and non-sequential read operations for a component over a specified time interval.  |
| Write I/O Rate (overall)              | Average number of I/O operations per second for both sequential and non-sequential write operations for a component over a specified time interval. |
| Total I/O Rate (overall)              | Average number of I/O operations per second for both sequential and nonsequential read and write operations.                                        |
| Disk to Cache Transfer Rate           | Average number of I/O operations (track transfers) per second for disk to cache transfers, for a particular component over an interval.             |
| Cache to Disk Transfer Rate           | Average number of I/O operations (track transfers) per second for cache to disk transfers, for a particular component over an interval.             |

| Column                               | Description                                                                                                                                                                                                |
|--------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Write-cache Delay Percentage         | Percentage of I/O operations that were delayed because of write-cache space constraints or other conditions, for a particular component over an interval. (The ratio of delayed operations to total I/Os.) |
| Write-cache Overflow Percentage      | Percentage of write operations that were delayed because of lack of write-cache space, for a particular component over an interval.                                                                        |
| Write-cache Flush-through Percentage | Percentage of write operations that were processed in flush-through write mode, for a particular component over an interval.                                                                               |
| Write Cache Write-through Percentage | Percentage of write operations that were processed in write-through write mode, for a particular component over an interval.                                                                               |
| Write-cache Delay I/O Rate           | Average number of I/O operations per second that were delayed because of write-cache space constraints or other conditions, for a particular component over an interval.                                   |
| Write-cache Overflow I/O Rate        | Average number of tracks per second that were delayed because of lack of write-cache space, for a particular component over an interval.                                                                   |
| Write Cache Flush-through I/O Rate   | Average number of tracks per second that were processed in flush-through write mode, for a particular component over an interval.                                                                          |
| Write Cache Write-through I/O Rate   | Average number of tracks per second that were processed in write-through write mode, for a particular component over an interval.                                                                          |
| Readahead Percentage of Cache Hits   | Percentage of all read cache hits which occurred on prestaged data.                                                                                                                                        |
| Dirty Write Percentage of Cache Hits | Percentage of all write cache hits which occurred on already dirty data in the cache.                                                                                                                      |

## Port Performance

This report shows performance data for ports. This report supports all storage systems.

| Column                         | Description                                                                                                           |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| Subsystem                      | ID of the storage system                                                                                              |
| Port                           | The ID of the specified port.                                                                                         |
| WWPN                           | The World Wide Port Name (WWPN) of the selected port.                                                                 |
| Time                           | Date and time that the data was collected.                                                                            |
| Interval                       | The actual duration of the sample interval (in seconds) for the performance data that was collected.                  |
| Port Send Bandwidth Percentage | The approximate bandwidth utilization percentage for send operations by a port based on its current negotiated speed. |

| Column                              | Description                                                                                                                                               |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port Receive Bandwidth Percentage   | The approximate bandwidth utilization percentage for receive operations by this port, based on its current negotiated speed.                              |
| Overall Port Bandwidth Percentage   | The approximate bandwidth utilization percentage for send and receive operations by this port.                                                            |
| Port Send Utilization Percentage    | The average amount of time that the port was busy sending data over a time interval.                                                                      |
| Port Receive Utilization Percentage | The average amount of time that the port was busy receiving data over a time interval.                                                                    |
| Overall Port Utilization Percentage | The average amount of time that the port was busy sending or receiving data over a time interval.                                                         |
| Port Send I/O Rate                  | The average number of input/output (I/O) operations per second for send operations for a port over a time interval.                                       |
| Port Receive I/O Rate               | Average number of I/O operations per second of receive operations for a port over a time interval.                                                        |
| Total Port I/O Rate                 | Average number of I/O operations per second for send and receive operations for a port over a time interval.                                              |
| Port Send Data Rate                 | The average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for send (read) operations for a port over a time interval.      |
| Port Receive Data Rate              | The average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for receive (write) operations for a port over a time interval.  |
| Total Port Data Rate                | The average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for send and receive operations for a port over a time interval. |
| Port to Host Send I/O Rate          | The average number of exchanges (I/Os) per second sent to host computers by a component over a time interval.                                             |
| Port to Host Receive I/O Rate       | The average number of exchanges (I/Os) per second received from host computers by a component over a time interval.                                       |
| Total Port to Host I/O Rate         | The average number of exchanges (I/Os) per second transmitted between host computers and a component over a time interval.                                |
| Port to Disk Send I/O Rate          | The average number of exchanges (I/Os) per second sent to storage subsystems by a component over a time interval.                                         |
| Port to Disk Receive I/O Rate       | The average number of exchanges (I/Os) per second received from storage subsystems by a component over a time interval.                                   |

| Column                         | Description                                                                                                                                 |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| Total Port to Disk I/O Rate    | The average number of exchanges (I/Os) per second transmitted between storage subsystems and a component over a time interval.              |
| Port to Host Send Data Rate    | The average number of megabytes (2 <sup>20</sup> bytes) per second sent to host computers by a component over a time interval.              |
| Port to Host Receive Data Rate | The average number of megabytes (2 <sup>20</sup> bytes) per second received from host computers by a component over a time interval.        |
| Total Port to Host Data Rate   | The average number of megabytes (2 <sup>20</sup> bytes) per second transmitted between host computers and a component over a time interval. |
| Port to Disk Send Data Rate    | Average number of megabytes (2 <sup>20</sup> bytes) per second sent to storage subsystems by a component over a time interval.              |
| Port to Disk Receive Data Rate | Average number of megabytes (2 <sup>20</sup> bytes) per second received from storage subsystems by a component over a time interval.        |
| Total Port Data Rate           | Average number of megabytes (2 <sup>20</sup> bytes) per second transmitted between storage subsystems and a component over a time interval. |
| Port Send Response Time        | The average number of milliseconds that it took to service each send (read) operation for a port over a time interval.                      |
| Port Receive Response Time     | The average number of milliseconds that it took to service each receive (write) operation for a port over a time interval.                  |
| Total Port Response Time       | The average number of milliseconds that it took to service each operation (send and receive) for a port over a time interval.               |
| Port Send Transfer Size        | The average number of KB sent per I/O by a port over a time interval.                                                                       |
| Port Receive Transfer Size     | The average number of KB received per I/O by a port over a time interval.                                                                   |
| Total Port Transfer Size       | The average number of KB transferred per I/O by a port over a time interval.                                                                |

## Subsystem Performance

This report shows performance data for storage systems. This report supports all storage systems.

| Column    | Description                                |
|-----------|--------------------------------------------|
| Subsystem | ID of the storage system.                  |
| Time      | Date and time that the data was collected. |

| Column                   | Description                                                                                                                                                        |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interval                 | The actual duration of the sample interval (in seconds) for the performance data that was collected.                                                               |
| Read I/O Rate (overall)  | Average number of I/O operations per second for both sequential and non-sequential read operations for a component over a specified time interval.                 |
| Write I/O Rate (overall) | Average number of I/O operations per second for both sequential and non-sequential write operations for a component over a specified time interval.                |
| Total I/O Rate (overall) | Average number of I/O operations per second for both sequential and nonsequential read and write operations.                                                       |
| Read Data Rate           | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for read operations for a component over a specified time interval.           |
| Write Data Rate          | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for write operations for a component over a specified time interval.          |
| Total Data Rate          | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for read and write operations for a component over a specified time interval. |
| Read Response Time       | Average number of milliseconds that it took to service each read operation for a component over a specified time interval.                                         |
| Write Response Time      | Average number of milliseconds that it took to service each write operation for a component over a specified time interval.                                        |
| Overall Response Time    | Average number of milliseconds that it took to service each I/O operation (read and write) for a component over a specified time interval.                         |

### Top Active Volumes Cache Hit Performance

This report summarizes the primary metrics for active volumes, prioritized by cache hit ratio. Only volumes with an input/output (I/O) rate that is greater than 0 will be included in this report. If all volumes are inactive, this report will be empty. This report supports all storage systems.

| Column    | Description                                                                                          |
|-----------|------------------------------------------------------------------------------------------------------|
| Subsystem | ID of the storage system.                                                                            |
| Volume    | The ID of the volume.                                                                                |
| Time      | Date and time that the data was collected.                                                           |
| Interval  | The actual duration of the sample interval (in seconds) for the performance data that was collected. |



| Column                                | Description                                                                                                                                                        |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Read Cache Hits Percentage (overall)  | Overall read cache percentage for both sequential and nonsequential read operations.                                                                               |
| Write Cache Hits Percentage (overall) | Percentage of cache hits for both sequential and nonsequential write operations.                                                                                   |
| Total Cache Hits Percentage (overall) | Percentage of cache hits for both sequential and nonsequential read and write operations.                                                                          |
| Readahead Percentage of Cache Hits    | Percentage of all read cache hits which occurred on prestaged data.                                                                                                |
| Dirty Write Percentage of Cache Hits  | Percentage of all write cache hits which occurred on already dirty data in the cache.                                                                              |
| Read I/O Rate (overall)               | Average number of I/O operations per second for both sequential and non-sequential read operations for a component over a specified time interval.                 |
| Write I/O Rate (overall)              | Average number of I/O operations per second for both sequential and non-sequential write operations for a component over a specified time interval.                |
| Total I/O Rate (overall)              | Average number of I/O operations per second for both sequential and nonsequential read and write operations.                                                       |
| Read Data Rate                        | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for read operations for a component over a specified time interval.           |
| Write Data Rate                       | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for write operations for a component over a specified time interval.          |
| Total Data Rate                       | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for read and write operations for a component over a specified time interval. |
| Read Response Time                    | Average number of milliseconds that it took to service each read operation for a component over a specified time interval.                                         |
| Write Response Time                   | Average number of milliseconds that it took to service each write operation for a component over a specified time interval.                                        |
| Overall Response Time                 | Average number of milliseconds that it took to service each I/O operation (read and write) for a component over a specified time interval.                         |

## Top Volumes Data Rate Performance

This report summarizes the primary performance metrics for volumes, prioritized by data rate. This report supports all storage systems.

| Column                   | Description                                                                                                                                                        |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Subsystem                | ID of the storage system                                                                                                                                           |
| Volume                   | The ID of the specified volume.                                                                                                                                    |
| Time                     | Date and time that the data was collected.                                                                                                                         |
| Interval                 | The actual duration of the sample interval (in seconds) for the performance data that was collected.                                                               |
| Read Data Rate           | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for read operations for a component over a specified time interval.           |
| Write Data Rate          | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for write operations for a component over a specified time interval.          |
| Total Data Rate          | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for read and write operations for a component over a specified time interval. |
| Read I/O Rate (overall)  | Average number of I/O operations per second for both sequential and non-sequential read operations for a component over a specified time interval.                 |
| Write I/O Rate (overall) | Average number of I/O operations per second for both sequential and non-sequential write operations for a component over a specified time interval.                |
| Total I/O Rate (overall) | Average number of I/O operations per second for both sequential and nonsequential read and write operations.                                                       |
| Read Response Time       | Average number of milliseconds that it took to service each read operation for a component over a specified time interval.                                         |
| Write Response Time      | Average number of milliseconds that it took to service each write operation for a component over a specified time interval.                                        |
| Overall Response Time    | Average number of milliseconds that it took to service each I/O operation (read and write) for a component over a specified time interval.                         |

### Top Volumes Disk Performance

This report summarizes the primary performance metrics for volumes, prioritized by disk rate. This report supports all storage systems.

| Column    | Description                                |
|-----------|--------------------------------------------|
| Subsystem | ID of the storage system.                  |
| Volume    | The ID of the volume.                      |
| Time      | Date and time that the data was collected. |

| <b>Column</b>                        | <b>Description</b>                                                                                                                                                                                         |
|--------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interval                             | The actual duration of the sample interval (in seconds) for the performance data that was collected.                                                                                                       |
| Disk to Cache Transfer Rate          | Average number of I/O operations (track transfers) per second for disk to cache transfers, for a particular component over an interval.                                                                    |
| Cache to Disk Transfer Rate          | Average number of I/O operations (track transfers) per second for cache to disk transfers, for a particular component over an interval.                                                                    |
| Write-cache Delay Percentage         | Percentage of I/O operations that were delayed because of write-cache space constraints or other conditions, for a particular component over an interval. (The ratio of delayed operations to total I/Os.) |
| Write-cache Overflow Percentage      | Percentage of write operations that were delayed because of lack of write-cache space, for a particular component over an interval.                                                                        |
| Write-cache Flush-through Percentage | Percentage of write operations that were processed in flush-through write mode, for a particular component over an interval.                                                                               |
| Write Cache Write-through Percentage | Percentage of write operations that were processed in write-through write mode, for a particular component over an interval.                                                                               |
| Write-cache Delay I/O Rate           | Average number of I/O operations per second that were delayed because of write-cache space constraints or other conditions, for a particular component over an interval.                                   |
| Write-cache Overflow I/O Rate        | Average number of tracks per second that were delayed because of lack of write-cache space, for a particular component over an interval.                                                                   |
| Write Cache Flush-through I/O Rate   | Average number of tracks per second that were processed in flush-through write mode, for a particular component over an interval.                                                                          |
| Write Cache Write-through I/O Rate   | Average number of tracks per second that were processed in write-through write mode, for a particular component over an interval.                                                                          |
| Read I/O Rate (overall)              | Average number of I/O operations per second for both sequential and non-sequential read operations for a component over a specified time interval.                                                         |
| Write I/O Rate (overall)             | Average number of I/O operations per second for both sequential and non-sequential write operations for a component over a specified time interval.                                                        |
| Total I/O Rate (overall)             | Average number of I/O operations per second for both sequential and nonsequential read and write operations.                                                                                               |

| Column                | Description                                                                                                                                                        |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Read Data Rate        | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for read operations for a component over a specified time interval.           |
| Write Data Rate       | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for write operations for a component over a specified time interval.          |
| Total Data Rate       | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for read and write operations for a component over a specified time interval. |
| Read Response Time    | Average number of milliseconds that it took to service each read operation for a component over a specified time interval.                                         |
| Write Response Time   | Average number of milliseconds that it took to service each write operation for a component over a specified time interval.                                        |
| Overall Response Time | Average number of milliseconds that it took to service each I/O operation (read and write) for a component over a specified time interval.                         |

### Top Volumes I/O Rate Performance

This report summarizes the primary performance metrics for volumes, prioritized by I/O rate. This report supports all storage systems.

| Column                   | Description                                                                                                                                              |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Subsystem                | ID of the storage system.                                                                                                                                |
| Volume                   | The ID of the specified volume.                                                                                                                          |
| Time                     | Date and time that the data was collected.                                                                                                               |
| Interval                 | The actual duration of the sample interval (in seconds) for the performance data that was collected.                                                     |
| Read I/O Rate (overall)  | Average number of I/O operations per second for both sequential and non-sequential read operations for a component over a specified time interval.       |
| Write I/O Rate (overall) | Average number of I/O operations per second for both sequential and non-sequential write operations for a component over a specified time interval.      |
| Total I/O Rate (overall) | Average number of I/O operations per second for both sequential and nonsequential read and write operations.                                             |
| Read Data Rate           | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for read operations for a component over a specified time interval. |

| Column                | Description                                                                                                                                                        |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Write Data Rate       | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for write operations for a component over a specified time interval.          |
| Total Data Rate       | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for read and write operations for a component over a specified time interval. |
| Read Response Time    | Average number of milliseconds that it took to service each read operation for a component over a specified time interval.                                         |
| Write Response Time   | Average number of milliseconds that it took to service each write operation for a component over a specified time interval.                                        |
| Overall Response Time | Average number of milliseconds that it took to service each I/O operation (read and write) for a component over a specified time interval.                         |

### Top Volumes Response Performance

This report summarizes the primary performance metrics for volumes, prioritized by response rate. This report supports all storage systems.

| Column                   | Description                                                                                                                                        |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| Subsystem                | ID of the storage system.                                                                                                                          |
| Volume                   | The ID of the volume.                                                                                                                              |
| Time                     | Date and time that the data was collected.                                                                                                         |
| Interval                 | The actual duration of the sample interval (in seconds) for the performance data that was collected.                                               |
| Read Response Time       | Average number of milliseconds that it took to service each read operation for a component over a specified time interval.                         |
| Write Response Time      | Average number of milliseconds that it took to service each write operation for a component over a specified time interval.                        |
| Overall Response Time    | Average number of milliseconds that it took to service each I/O operation (read and write) for a component over a specified time interval.         |
| Peak Read Response Time  | The peak (worst) response time for disk reads for a component over a time interval.                                                                |
| Peak Write Response Time | The peak (worst) response time for disk writes for a component over a time interval.                                                               |
| Read I/O Rate (overall)  | Average number of I/O operations per second for both sequential and non-sequential read operations for a component over a specified time interval. |

| Column                   | Description                                                                                                                                                        |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Write I/O Rate (overall) | Average number of I/O operations per second for both sequential and non-sequential write operations for a component over a specified time interval.                |
| Total I/O Rate (overall) | Average number of I/O operations per second for both sequential and nonsequential read and write operations.                                                       |
| Read Data Rate           | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for read operations for a component over a specified time interval.           |
| Write Data Rate          | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for write operations for a component over a specified time interval.          |
| Total Data Rate          | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for read and write operations for a component over a specified time interval. |

## System reports - Fabric Manager

Use these predefined reports to view performance information about fabrics.

### Port Connections

Use the Port Connections report to view information about the ports on the switches in your storage area network (SAN). This report includes information about the port connections and status, as well as the port and switch IDs.

|                      |                                                                                                            |
|----------------------|------------------------------------------------------------------------------------------------------------|
| Fabric ID            | The ID of the fabric that the switch is in.                                                                |
| Switch               | The label value of the switch.                                                                             |
| Switch Port          | The ID of the port ( <i>domainID,port_number</i> ).                                                        |
| Slot                 | The number for the blade that the port is on.                                                              |
| Index                | The unique number for the port across the whole switch.                                                    |
| Switch Port State    | The state of the switch port: Normal, Warning, or Critical.                                                |
| Connection State     | The state of the connection between the two ports: Normal, Warning, or Critical.                           |
| Connected Port       | The port that the switch port is connected to. This value is the WWN of the device connected to that port. |
| Connected Port State | The state of the connected port: Normal, Warning, or Critical.                                             |
| Device               | The device that the connected port is on.                                                                  |
| Device State         | The state of the device that the connected port is on: Normal, Warning, or Critical.                       |

### SAN Assets (ALL)

Use the SAN Assets (All) report to generate and view reports that provide information about all the assets discovered by IBM Tivoli Storage Productivity Center.

The following table describes the columns on the SAN Assets (ALL) report. All the columns in the report are for display only.

|                   |                                                                        |
|-------------------|------------------------------------------------------------------------|
| <b>Fabric ID</b>  | The ID of the fabric that the asset is in.                             |
| <b>Type</b>       | The type of asset discovered.                                          |
| <b>Label</b>      | The label of the asset.                                                |
| <b>Identifier</b> | The unique identifier of the storage asset. For example, a switch WWN. |
| <b>Status</b>     | The state of the device, either Normal, Warning, or Critical.          |

### **SAN Assets (Connected Devices)**

Use the SAN Assets (Connected Devices) report to generate and view reports that provide information about all the connected devices discovered by IBM Tivoli Storage Productivity Center.

The following table describes the columns on the SAN Assets (Connected Devices) report. All the columns in the report are for display only.

|                   |                                                                             |
|-------------------|-----------------------------------------------------------------------------|
| <b>Fabric ID</b>  | The ID of the fabric that the asset is in.                                  |
| <b>Type</b>       | The type of asset that was discovered.                                      |
| <b>Label</b>      | The label of the asset.                                                     |
| <b>Identifier</b> | The unique identifier of the asset. For example, the device worldwide name. |
| <b>Status</b>     | The status of the asset: Normal, Warning, or Critical.                      |

### **SAN Assets (Switches)**

Use the SAN Assets (Switches) report to generate and view reports that provide information about switches in your SAN. These reports provide information about the switch ID, type, manufacturer, WWN, location, and other details.

The following table describes the columns on the SAN Assets (Switches) report. All the columns in the report are for display only.

|                            |                                                                                                                           |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Fabric ID</b>           | The ID of the fabric that the switch is in.                                                                               |
| <b>Type</b>                | The type is always "Switch".                                                                                              |
| <b>Switch Type</b>         | The type of switch discovered. For example, a physical switch.                                                            |
| <b>Label</b>               | The label of the switch.                                                                                                  |
| <b>Status</b>              | The status of the switch. The possible states are <b>Normal</b> and <b>Missing</b> .                                      |
| <b>IP Address</b>          | The IP address of the switch.                                                                                             |
| <b>Vendor ID</b>           | The manufacturer of the switch.                                                                                           |
| <b>Model</b>               | The model name and number of the switch.                                                                                  |
| <b>Version</b>             | The version information of the switch.                                                                                    |
| <b>Serial Number</b>       | The serial number of the switch.                                                                                          |
| <b>Object ID</b>           | The object ID of the switch.                                                                                              |
| <b>Management ID</b>       | The management ID of the switch.                                                                                          |
| <b>Management Address</b>  | The URL of the switch element manager.                                                                                    |
| <b>Domain</b>              | The ID of the domain that the switch is in.                                                                               |
| <b>Contact</b>             | The contact information for the switch.                                                                                   |
| <b>WWN</b>                 | The world-wide name (WWN) of the switch.                                                                                  |
| <b>Location</b>            | The physical location of the switch.                                                                                      |
| <b>Physical Switch WWN</b> | The WWN for the parent or physical switch of a virtual switch. If this is not a virtual switch then this column is empty. |

## Switch Performance

Use the Switch Performance report to generate and view reports that provide information about the free and occupied ports on the switches in your SAN. These reports provide information about the port connections and status, as well as the port and switch IDs, data and packet transfer rates, and other details. This report can only be generated for switches that have had performance monitors run on them.


This following table describes the columns on the Switch Performance report. All the columns in the report are for display only.

|                                    |                                                                            |
|------------------------------------|----------------------------------------------------------------------------|
| <b>Switch</b>                      | The ID of the switch.                                                      |
| <b>Time</b>                        | The date and time of the data collection.                                  |
| <b>Interval</b>                    | The size of the sample interval, in seconds.                               |
| <b>Port Send Packet Rate</b>       | The rate at which packets are sent through this port.                      |
| <b>Port Receive Packet Rate</b>    | The rate at which packets are received through this port.                  |
| <b>Total Port Packet Rate</b>      | The total rate at which packet operations are performed through this port. |
| <b>Port Send Data Rate</b>         | The rate at which data is sent through this port.                          |
| <b>Port Receive Data Rate</b>      | The rate at which data is received through this port.                      |
| <b>Total Port Data Rate</b>        | The total rate at which data is processed through this port.               |
| <b>Port Peak Send Data Rate</b>    | The peak rate at which data is sent through this switch.                   |
| <b>Port Peak Receive Data Rate</b> | The peak rate at which data is received through this switch.               |
| <b>Port Send Packet Size</b>       | The size of the data block that is sent out through this switch.           |
| <b>Port Receive Packet Size</b>    | The size of the data block that is received in through this switch.        |
| <b>Overall Port Packet Size</b>    | The total size of the data block that is transferred by the switch.        |
| <b>Error Frame Rate</b>            | The rate at which error frames are generated by the switch.                |
| <b>Dumped Frame Rate</b>           | The rate at which frames are dumped by the switch.                         |
| <b>Link Failure Rate</b>           | The rate at which failed links are generated by the switch.                |

## Switch Port Errors

Use the Switch Port Errors report to generate and view reports that provide information about the errors being generated by ports on the switches in your SAN. These reports provide information about the error frames, dumped frames, link failures, port and switch IDs, and other details. This report can only be generated for switches that have had performance monitors run on them.

To view more information for a specific file:

- Click  . Detailed information per computer will be displayed in a **Network** tab. The **Network** tab displays file summary by computer.

This table describes the fields on the Switch Port Errors report. All the fields on the report are for display only.

|                              |                                                                  |
|------------------------------|------------------------------------------------------------------|
| <b>Switch</b>                | The ID of the switch.                                            |
| <b>Port</b>                  | The ID of the port.                                              |
| <b>WWPN</b>                  | The World Wide Port Name (WWPN) of the selected port.            |
| <b>Slot</b>                  | The number for the blade that the port is on.                    |
| <b>Index</b>                 | The unique number for the port across the whole switch.          |
| <b>Time</b>                  | The date and time of the data collection.                        |
| <b>Interval</b>              | The size of the sample interval, in seconds.                     |
| <b>Port Send Packet Size</b> | The size of the data block that is sent out through this switch. |



|                                      |                                                                                                                                                                                                                                      |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Port Receive Packet Size</b>      | The size of the data block that is received in through this switch.                                                                                                                                                                  |
| <b>Overall Port Packet Size</b>      | The total size of the data block that is transferred by the switch                                                                                                                                                                   |
| <b>Error Frame Rate</b>              | The rate at which error frames are generated by the switch.                                                                                                                                                                          |
| <b>Dumped Frame Rate</b>             | The rate at which frames are dumped by the switch.                                                                                                                                                                                   |
| <b>Link Failure Rate</b>             | The rate at which failed links are generated by the switch.                                                                                                                                                                          |
| <b>Loss of Sync Rate</b>             | The average number of times per second that synchronization was lost after the last reset of the device, for a particular component over a particular time interval.                                                                 |
| <b>Loss of Signal Rate</b>           | The average number of times per second that the signal was lost after the last reset of the device, for a particular component over a particular time interval.                                                                      |
| <b>CRC Error Rate</b>                | The average number of frames received per second in which the CRC in the frame did not match the CRC computed by the receiver, for a particular component over a particular time interval.                                           |
| <b>Short Frame Rate</b>              | The average number of frames received per second that were shorter than 28 octets (24 header + 4 CRC), not including any start of frame/end of frame bytes, for a particular component over a particular time interval.              |
| <b>Long Frame Rate</b>               | The average number of frames received per second that were longer than 2140 octets (24 header + 4 CRC + 2112 data), not including any start of frame/end of frame bytes, for a particular component over a particular time interval. |
| <b>Encoding Disparity Error Rate</b> | The average number of disparity errors received per second, for a particular component over a particular time interval.                                                                                                              |

This table describes columns that are not part of the Switch Port Errors system report. However, you can include them, and save the resulting report in the `<user-id> Reports` node.


|                                    |                                                                            |
|------------------------------------|----------------------------------------------------------------------------|
| <b>Port Send Packet Rate</b>       | The rate at which packets are sent through this port.                      |
| <b>Port Receive Packet Rate</b>    | The rate at which packets are received through this port.                  |
| <b>Total Port Packet Rate</b>      | The total rate at which packet operations are performed through this port. |
| <b>Port Send Data Rate</b>         | The rate at which data is sent through this port.                          |
| <b>Port Receive Data Rate</b>      | The rate at which data is received through this port.                      |
| <b>Total Port Data Rate</b>        | The total rate at which data is processed through this port.               |
| <b>Port Peak Send Data Rate</b>    | The peak rate at which data is sent through this switch.                   |
| <b>Port Peak Receive Data Rate</b> | The peak rate at which data is received through this switch.               |

## Top Switch Ports Data Rate Performance

Use the Top Switch Port Data Rate Performance report to generate and view reports that provide information about the switches in your SAN that have the top data rates. This report can only be generated for switches that have had performance monitors run on them.

This report only displays rows for the 25 highest data rates from the switch ports monitored. It displays those rows in descending order by Total Port Data Rate. You can adjust the number of rows shown on the **Selection** tab.

**To view more information for a specific file:**

- Click  . Detailed information per computer will be displayed in a **Network** tab. The **Network** tab displays the file summary by computer.

This section describes the fields on the Top Switch Ports Data Rate Performance report. All the fields on the report are for display only.

|                               |                                                                                |
|-------------------------------|--------------------------------------------------------------------------------|
| <b>Switch</b>                 | The ID of the switch.                                                          |
| <b>Port</b>                   | The ID of the port.                                                            |
| <b>Time</b>                   | The date and time of the data collection.                                      |
| <b>Interval</b>               | The size of the sample interval, in seconds.                                   |
| <b>Total Port Data Rate</b>   | The total rate at which data is processed through this port.                   |
| <b>Total Port Packet Rate</b> | The total rate at which the packet operations are performed through this port. |

This table describes columns that are not part of the Top Switch Ports Data Rate Performance report. However, you can include them, and save the resulting report in the *<user> Reports* node.


|                                    |                                                                     |
|------------------------------------|---------------------------------------------------------------------|
| <b>Port Send Packet Rate</b>       | The rate at which packets are sent through this port.               |
| <b>Port Receive Packet Rate</b>    | The rate at which packets are received through this port.           |
| <b>Port Send Data Rate</b>         | The rate at which data is sent through this port.                   |
| <b>Port Receive Data Rate</b>      | The rate at which data is received through this port.               |
| <b>Port Peak Send Data Rate</b>    | The peak rate at which data is sent through this switch.            |
| <b>Port Peak Receive Data Rate</b> | The peak rate at which data is received through this switch.        |
| <b>Port Send Packet Size</b>       | The size of the data block that is sent out through this switch.    |
| <b>Port Receive Packet Size</b>    | The size of the data block that is received in through this switch. |
| <b>Overall Port Packet Size</b>    | The total size of the data block that is transferred by the switch  |
| <b>Error Frame Rate</b>            | The rate at which error frames are generated by the switch.         |
| <b>Dumped Frame Rate</b>           | The rate at which frames are dumped by the switch.                  |
| <b>Link Failure Rate</b>           | The rate at which failed links are generated by the switch.         |
| <b>WWPN</b>                        | The World Wide Name identifier for the FC port.                     |
| <b>Slot</b>                        | The number for the blade that the port is on.                       |
| <b>Index</b>                       | The unique number for the port across the whole switch.             |

## Top Switch Ports Packet Rate Performance

Use the Top Switch Ports Packet Rate Performance report to generate and view reports that provide information about the switches in your SAN that have the top packet rates. This report can only be generated for switches that have had performance monitors run on them.

This report only displays rows for the 25 highest packet rates from the switch ports monitored. It displays those rows in descending order by Total Port Packet Rate. You can adjust the number of rows shown on the Selection tab.

### To view more information for a specific file:

- Click  . Detailed information per computer will be displayed in a **Network** tab. The **Network** tab displays file summary by computer.

This section describes the fields on the Top Switch Ports Packet Rate Performance report. All the fields on the report are for display only.

|               |                       |
|---------------|-----------------------|
| <b>Switch</b> | The ID of the switch. |
|---------------|-----------------------|

|                               |                                                                                |
|-------------------------------|--------------------------------------------------------------------------------|
| <b>Port</b>                   | The ID of the port.                                                            |
| <b>Time</b>                   | The date and time of the data collection.                                      |
| <b>Interval</b>               | The size of the sample interval, in seconds.                                   |
| <b>Total Port Packet Rate</b> | The total rate at which the packet operations are performed through this port. |
| <b>Total Port Data Rate</b>   | The total rate at which data is processed through this port.                   |

This table describes columns that are not part of the Top Switch Ports Packet Rate Performance report. However, you can include them, and save the resulting report in the *<user> Reports* node.

|                                    |                                                                     |
|------------------------------------|---------------------------------------------------------------------|
| <b>Port Send Packet Rate</b>       | The rate at which packets are sent through this port.               |
| <b>Port Receive Packet Rate</b>    | The rate at which packets are received through this port.           |
| <b>Port Send Data Rate</b>         | The rate at which data is sent through this port.                   |
| <b>Port Receive Data Rate</b>      | The rate at which data is received through this port.               |
| <b>Port Peak Send Data Rate</b>    | The peak rate at which data is sent through this switch.            |
| <b>Port Peak Receive Data Rate</b> | The peak rate at which data is received through this switch.        |
| <b>Port Send Packet Size</b>       | The size of the data block that is sent out through this switch.    |
| <b>Port Receive Packet Size</b>    | The size of the data block that is received in through this switch. |
| <b>Overall Port Packet Size</b>    | The total size of the data block that is transferred by the switch  |
| <b>Error Frame Rate</b>            | The rate at which error frames are generated by the switch.         |
| <b>Dumped Frame Rate</b>           | The rate at which frames are dumped by the switch.                  |
| <b>Link Failure Rate</b>           | The rate at which failed links are generated by the switch.         |
| <b>WWPN</b>                        | The World Wide Name identifier for the FC port.                     |
| <b>Slot</b>                        | The number for the blade that the port is on.                       |
| <b>Index</b>                       | The unique number for the port across the whole switch.             |

## Batch reports

Use the information in this section to learn how to generate and view batch reports. These reports are available in the IBM Tivoli Storage Productivity Center > **Reporting** > **My Reports** > **Batch Reports** node in the navigation tree.

A batch report represents any Tivoli Storage Productivity Center report that you want to run on a regular basis and save its data to a file. Schedule batch reports to run immediately, run once at a specified date and time, or run repeatedly at a specified date, time, cycle, and day. You can also select the time zone to use when running the report.

You can view the information in a generated batch report file directly or import it into an application of your choice. Batch reports provide a convenient and powerful way for you to save report definitions and schedule when to run those reports.



### Creating batch reports

Create a batch report.

To create a batch report follow this procedure:

1. Expand **IBM Tivoli Storage Productivity Center** → **Reporting** → **My Reports**.
2. Right-click the **Batch Reports** node and select **Create Batch Report**.

**Note:** If the batch report you want to create is similar to an existing one, click the name of the existing report, edit it and save it using a different name.

3. Click the **Report** tab and select the type of report you want to create and move it from **Available** to **Current Selections** by clicking . See "Selecting batch report options" for details of the tab contents.
4. Click the **Selection** tab and refine the criteria for selecting and displaying report data. Some options might not be available for all reports. Use the **Selection** page to:
  - Select a profile for controlling which data is collected. This option is not available for all reports. See "Using profiles" on page 352 for information about profiles.
  - Click **Selection** to specify the resources (for example, computers) that will be displayed in the report. This option is not available for all reports. For details about selecting resources, see "Selecting resources to appear on a report" on page 796.
  - Click **Filters** to apply filters the data that you want to display (for example, show only the machines that contain over 80% used space). For details about applying filters, see "Filtering the resources that appear on a report" on page 799.
  - Select the columns you want to include in the report and move them from **Available Columns** to **Included Columns** by clicking .

**Note:**

- The values you set in **Available Columns** and **Included Columns** are not used for the following types of batch reports: History CSV File, PDF Chart, and HTML Chart.
  - The value you set in the **Summation Level** field is not used for the following types of batch reports: History CSV File, PDF Chart, and HTML Chart.
5. Click the **Options** tab of the Batch Report job window, specify the following:
    - The machine on which the report file will be saved
    - The report format (for example, HTML)
    - Whether to run a script when the report process is complete
    - The format for the name of the batch report

See "Selecting batch report options" for details of the tab contents.

6. In the **When to Run** tab of the Batch Report job window, specify when and how often to run the batch report. See "Scheduling a batch report" on page 1027 for details of the tab contents.
7. In the **Alert** tab of the Batch Report job window, define the conditions that will trigger an alert and the method of notification (for example, email). See "Defining alerts for batch reports" on page 1028 for details of the tab contents.
8. Click **File** → **Save** and type a name for the batch report. The batch report will be saved using your user ID as a prefix. For example, if you are logged onto Tivoli Storage Productivity Center as "Smith," the name of the batch report will be `Smith.report_name`.

### Selecting batch report options

Define the agent computer that you want to run the batch report on and specify the format of the file to be saved.

Use the **Options** page to define where the batch report is run and specify information about the output file.

1. Use the following fields to specify the location where a batch report is saved.

### TPC Server

Save the batch report on the computer where the Tivoli Storage Productivity Center server is located. Use the **Report Location** field to specify the directory where the batch report file is saved.

**Agent** Save the batch report on the computer where an agent is located. Use the **Report Location** field to specify the directory where the batch report file is saved.

You can save batch reports to computers on which the following agents are deployed:

- Data agents
- Storage Resource agents that are v4.2 or higher.

The **Select Agent** field becomes active if you select to save a batch report to an agent.

### Select Agent

Select the name of the computer on which an agent is deployed. This menu only displays computers on which the following agents are deployed:

- Data agents
- Storage Resource agents that are v4.2 or higher.

To save a batch report to a computer on which v4.1 of a Storage Resource agent is deployed, you must upgrade that agent to v4.2 or higher.

### Note:

- The type of report that is selected for the batch report determines the agents that are available in this list. Agents that do not support the selected report are not displayed. For example, agents for versions of the product before v4.1 are not available for Data Source reports.
- If the agent you select from this menu is removed from the computer before the batch report is run, an error occurs and the batch report is not generated. If this error occurs, edit the definition of the batch report to specify an existing agent and then run the report again.
- There are differences in the batch reports depending on whether a Storage Resource agent or a Data agent is selected. If you are using a Storage Resource agent to collect data, these additional columns are available in the batch reports using the Selection tab:
  - Is Encryptable
  - Is Encrypted

### Report Location

Specify the directory where you want to save the batch report. This field is only available for batch reports that are sent to the following locations:

- TPC Server
- Storage Resource agents that are v4.2 or higher.

For Data agents, batch reports are automatically saved to the following default locations:

### Windows

C:\Program Files\IBM\TPC\ca\subagents\TPC\Data\log\  
*computer\_name*\reports\

**Note:** If you upgraded to Tivoli Storage Productivity Center v4.2 or higher from a version before 3.1, the default location for batch reports is:

```
C:\Program Files\IBM\TPC\ep\subagents\TPC\Data\log\
computer_name\reports\
```

**UNIX or Linux**

```
/opt/IBM/TPC/ca/subagents/TPC/Data/log/computer_name/reports/
```

**Note:** If you upgraded to Tivoli Storage Productivity Center v4.2 or higher from a version before 3.1, the default location for batch reports is:

```
/opt/tivoli/ep/subagents/TPC/Data/log/computer_name/reports/
```

For Storage Resource agent batch reports, when you select a Storage Resource agent from the **Select Agent** drop-down list, you see the location for the batch reports.

**Windows**

```
SRA_location\agent\log\computername\reports
```

**UNIX or Linux**

```
SRA_location/agent/log/computername/reports
```

**Note:**

You can modify the location of the batch reports for Storage Resource agents and the TPC server only. You cannot modify the location of the batch reports for the Data agents.

You also have the option of saving the report on the Tivoli Storage Productivity Center server. The default path for the Tivoli Storage Productivity Center server is:

**Windows**

```
C:\Program Files\IBM\TPC\data\log\reports
```

**UNIX or Linux**

```
/opt/IBM/TPC/data/log/reports
```

2. In the **Report Type Specification** section, select the type of batch report you want to generate. You are limited to one type of output per batch report job and can select from the following output formats:

**CSV File**

Create a CSV (comma delimited) file of the batch report. Select **Include Headers** to include column headers in the batch report output file. Select **Include Totals** to include the Totals column in the report output file.

**Formatted File**

Create a formatted text file of the batch report.

**HTML File**

Create an HTML file of the batch report.

**History CSV File**

For reports with historical trending, you can export the historical data to a CSV file. The CSV file has the following format: *<object name, date, amount, date, amount>* and so on, where *<date, amount>* corresponds to each data point on the history chart.

## PDF Chart

Create a chart for the batch report in PDF format. After you select this option, you can determine the type of chart you want to create. The list of available chart types vary by the report you selected on the **Report** tab:

- Chart space usage
- History chart: File Count
- History chart: Free Space
- History chart: Free Space %
- History chart: Used Space
- History chart: Used Space %
- History chart: Capacity

### Note:

- You must use Adobe Acrobat Reader to view the resulting PDF. You can download Adobe Acrobat Reader from: <http://get.adobe.com/reader/>.
- This option is available for reports that support charting only.
- For agents that run on Linux and UNIX computers, you must perform additional configuration steps to enable this option. For information about how to configure the agent system, go to the Tivoli Storage Productivity Center Information Center. Search for the topic "Configuring IBM Tivoli Storage Productivity Center for batch reports on UNIX or Linux."
- The batch chart uses any customization settings that are defined for the corresponding "online" chart. "Online" charts are those charts generated using the Reporting facility.

For example, if you choose to generate a **TPC-wide Storage Space > File System Space > By Filesystem > PDF Chart > Chart space usage** batch report chart, any customization settings defined in the Reporting facility for **Data Manager > Reporting > TPC-wide Storage Space > File System Space > By Filesystem > Chart space usage for selected** is applied to the batch report chart.

Any batch chart with an associated online report that provides a choice between "selected" and "all" for its generated charts use the customization settings defined for the "selected" chart.

- If you select **PDF Chart > Pie Chart** from the Chart Customization window and one of the elements represented in the pie chart is much smaller than the other elements, the pie chart displays a graphical error. For example, if the pie chart for a batch report shows disk space, and the disk space used is 4 KB and the free space is 10 GB, a graphical error is displayed.

## HTML Chart

Select this option if you want to create a chart for the batch report in HTML format. When HTML Chart is selected, the drop-down list to the right becomes active. Use this drop-down list to select the type of chart you want to create. The list of available chart types vary by the report you selected on the **Report** tab:

- Chart space usage
- History chart: File Count
- History chart: Free Space

- History chart: Free Space %
- History chart: Used Space
- History chart: Used Space %
- History chart: Capacity

**Note:**

- If you generate charts for reports and these charts contain long path names or a large number of segments, the text on the charts might overlap and make them difficult to read.
- For agents that run on Linux and UNIX computers, you must perform additional configuration steps to enable this option. For information about how to configure the agent system, go to the Tivoli Storage Productivity Center Information Center. Search for the topic "Configuring IBM Tivoli Storage Productivity Center for batch reports on UNIX or Linux."
- The batch chart you generate uses any customization settings that are defined for the corresponding "online" chart. "Online" charts are those generated using the Reporting facility.

For example, if you choose to generate a **TPC-wide Storage Space > File System Space > By Filesystem > PDF Chart > Chart space usage** batch report chart, any customization settings defined in the Reporting facility for **Data Manager > Reporting > TPC-wide Storage Space > File System Space > By Filesystem > Chart space usage for selected** is applied to the batch report chart.

Any batch chart with an associated online report that provides a choice between "selected" and "all" for its generated charts use the customization settings defined for the "selected" chart.

**Customize this chart**

Click this button to use the Chart Customization window to customize the chart that is generated for the batch report. You will be able to select a chart type and chart orientation. If there is a customized chart for the corresponding online report, the customized chart is used for this batch report. However, you can override that customized chart by using this function.

The button is active only when **PDF Chart** or **HTML Chart** is selected.

3. In the **Classic Column Names Specification** section, check **Use Classic Column Names** to use the original names of report columns in the output for batch reports. If you do not select this option, the new names of columns are used. For your reference, the online help displays the original names of those columns within parentheses ( ) next to their new names.

**Note:** The titles of some reports have been updated. These new titles are always used for batch reports

4. In the **Script** section, check **Run Script** to run a script when the batch report is completed.

Select this option and enter the name of the script you want to run against the results of the batch report. For example, if you create a batch report for the most obsolete files on monitored computers, you can indicate a script that runs a program to archive all the files that appear in that report.

5. In the **Output File Specification** section, define the output file naming format for the batch report file.



### Enter Mask that will be used to create the output file name

Displays the parameters that determine the name of the output file. For example, if you use the default parameters, {Report creator} {Report name} {Report run number}, Tivoli Storage Productivity Center retrieves the values for these parameters and uses them to name the output file. Click any of the following parameters to include them in the mask for the name of an output file:

- Report creator (default)
- Report name (default)
- Report run number (default)
- Report title
- Report run date
- Report run time
- Unique number

**Note:** If you include {Report name} in the mask, Tivoli Storage Productivity Center uses the name of the batch report definition in the name of the output file. If you included NLS characters in the name of the batch report definition, errors might occur if you selected a Storage Resource agent as a destination for the batch report. You can use NLS characters in the name of a batch report definition in the following situations:

- TPC server or Data agent is selected as the destination for a batch report.
- A Storage Resource agent is selected as the destination for a batch report and the {Report name} parameter is not included in the mask for the output file name. If {Report name} is not used in the mask, it is recommended that other parameters are included in the mask to ensure that the output file name is unique. For example, {Report Creator} {Report Run Number} {Report Title} {Unique number}.

### Scheduling a batch report

Determine when and how often to run a batch report.

Use the **When to Run** page to specify when and how often to run the batch reporting job.

You have the following options for determining when to generate a report:

- Generate the report immediately
- Generate the report once at a specified time and date
- Generate the report repeatedly according to a schedule you specify

Select the option you want and then edit the time and date information accordingly. You can also use the fields in the How to handle time zones section to indicate the time zone that IBM Tivoli Storage Productivity Center should use when scheduling and running the batch report. For example, if you selected to run the job at 9:00 AM, you can use the fields in this section to run the batch report at:

- 9:00 AM local time of the agent running the batch report
- 9:00 AM of the time zone where the server resides
- 9:00 AM of a specific time zone, such as 9:00 AM CST or 9:00 GMT. For example, if you select 9:00 AM CST, the batch report will run at 7:00 AM on an agent located in the PST time zone.

**Local time in each timezone.** Select this option to use the time zone of the computer where the agent that collects data for the job is deployed. For probes of storage systems, tape libraries, and fabrics, the time zone of the Tivoli Storage Productivity Center server is used.

**Same Global time across all timezones.** Use the options in this section to apply the same global time across all the time zones where the probe is being run.

- **Use the timezone that the server runs in.** Select this option to use the time zone of the location where the Tivoli Storage Productivity Center server resides.
- **Use this timezone.** Select this option to indicate a specific time zone for when to run the job. When you select this option, the list box to the right of the field will activate and enable you to select from a wide range of global time zones.

## Defining alerts for batch reports

Define the alert that will be triggered if the batch report fails to run.

Use the **Alert** tab to define an alert that will be triggered if the report does not run on schedule.

To define alert values for a batch report, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center > Reporting -> My Reports** and right-click **Batch Reports**.
2. Select **Create Batch Reports**. The Create Batch Reports window opens.
3. For information on creating the batch report, see “Batch reports” on page 1021.
4. Click the **Alerts** tab.
5. Define the conditions that will trigger the alert. The default setting is **Report Failed**. No other alerting conditions are currently available.
6. Select alert mechanisms in addition to writing an error to the error log:
  - If you want to use Simple Network Management Protocol (SNMP) traps, IBM Tivoli Enterprise Console events, or email as an alert notification, you must specify disposition values in the Alert Dispositions window. Expand **Administrative Services > Configuration > Alert Disposition**.
  - If you want a script to be run after an event has been triggered, you must create the script before you create the alert.
7. Save the file:
  - a. Click **File > Save**.
  - b. The **Save As** window opens. Enter a name for the batch report and click **OK**.

### Editing alert values for a batch report

1. Expand **IBM Tivoli Storage Productivity Center > My Reports > Batch Reports**. The list of existing batch reports is displayed.
2. Click the batch report. The **Edit Batch Report** window opens.
3. In the **Alerts** tab, make any necessary changes.
4. Click **File > Save**.

### Related topics:

“Batch reports” on page 1021

## Creating batch reports for storage subsystem or switch performance

This topic describes how to create a batch report for storage subsystem performance or switch performance.

To create a batch report for storage subsystem performance or switch performance, follow these steps:

1. Expand **IBM Tivoli Storage Productivity Center > Reporting > My Reports**.
2. Right-click the **Batch Reports** node and select **Create Batch Report**.
3. Click the **Report** tab and select the type of report you want to create and move it from **Available** to **Current Selections** by clicking the right arrow button (>>). Select an item from **Storage Subsystem Performance** or **Switch Performance**.
4. Click the **Selection** tab. The Report Filter Specifications window displays.

### Selection button

This determines the resources upon which you want to report, such as the storage subsystems, the fibre-channel ports, and so forth. The Select Resources window displays. The type of resources displayed for selection depends on the type of report you previously selected on the **Report** tab. For example, if you selected a By Storage Subsystem report, all displayed resources will be storage subsystems. If you clear the box for a resource on the Select Resources window, a warning icon displays on the Selection button.

### Filter button

This further defines the data that you want to include in the report. Filters enable you to apply general rules to the report. When you click this button, the Edit Filter window displays.

For example, to select only data associated with subsystems that have a name ending with "123-IBM", select the "Subsystem" column, the "LIKE" operator, and type "\*123-IBM" (without the quotation marks) in the Value 1 field. To select only data which has a Read I/O Rate greater than 0, select the "Read I/O Rate (overall)" column, the ">" operator, and type "0" (without the quotation marks) in the Value 1 field.

If a filter is defined for this report, a warning icon displays on the **Filter** button.

### Time range selectors

All performance reports display historical performance data. Use the time range selectors to specify the exact time range that you want the report to cover. Click one of the following radio buttons to select the type of historical report you want:

#### Display latest performance data

This report contains only the latest available performance data for the selected resources. Note that even the latest data available might still be somewhat old, depending on whether data was recently collected for the resources. Including the time column in the report will help you determine how old the latest available data actually is. Regardless of the age of the data, clicking this button ensures that only a single row (representing a single point in time) for each selected resource is included in the report. If no data has ever been collected for a particular resource, then that resource is not included in the report. If no data has ever been collected for any resource, then the resulting report will be empty.

### Display historic performance data using absolute time

This report contains only data that falls within a specific time range. Select the time range using the **From** and **To** selectors. The specified from and to values are interpreted as inclusive. The resulting report contains all the performance data that is available within the specified time range, for each selected resource. This means that the data can consist of potentially many rows per resource, each with a unique timestamp. If no data has been collected for a particular resource during the specified time range, that resource is not included in the report. If no data has been collected for any resource during the specified time range, then the resulting report will be empty.

To determine which historical performance records fall within the specified time range, the **beginning timestamp** of the collection intervals is used for comparison. If an historical record contains performance data for an interval of time from 1:00 PM to 1:05 PM, that record would be included in the report for a time range of 1:00 PM to 8:00 PM. However, that same record would not be included if a time range of 1:02 PM to 8:02 PM was specified, because it is the beginning of the time period (for example, 1:00 PM in this example) that is used to determine inclusion in the specified time interval.

Specifying an absolute time range is not suitable for batch reports that are run on a repeating schedule, because the specified time range is fixed and will not automatically be updated based on the most current date and time. The result will be that the same report is generated every time that the batch report is run. Instead, use a relative time range for this type of report.

### Display historic performance data using relative time

This report contains only data that falls within a specific time range extending backward from the current date and time. Specify the number of days in the desired time range, using the **days ago until now** field. The report will be exactly the same as a report using absolute time, with the **To** selector set to the current date and time, and the **From** selector set to the current date and time minus the number of **days ago until now**. However, when regenerating this report periodically, as the current date and time advances, the data included in the report will advance as well. As a result, this is the suitable option to select when defining batch reports that are run on a repeated schedule.

Note that the same criteria applies as if determining which historical performance records fall within the specified time range. Only the beginning timestamp of the collection intervals is used for comparison. This means that data for the current hour or the current day is not included in the report when selecting hourly or daily summation data, because data for the current hour or current day is only saved in the database when the hour or day is complete. As an alternative, consider using the first radio button, **Display latest performance data**, to see a report showing the most recent daily summation data.

## Summation Level

Use this selector to select which type of performance data records to include in the report. The Tivoli Storage Productivity Center Performance Manager collects and stores three different types of performance data:

### By sample

Represents the most detailed data. There is usually one sample record per resource for every **n** minutes, where **n** is the performance monitor interval length that was specified when the monitor was defined. Each sample represents the average performance of the resource over those **n** minutes. A sample record with a timestamp of **t** will be saved at time **t+n**. In other words, the beginning of each **n** minute time interval is used for display and for time range matching.

### Hourly

Represents an hourly average of the performance of the specified resources. There will be one hourly record written every hour, per resource. Each such record will be written at or shortly after each hour mark on the (device) clock. For example, a record with a timestamp of 12:00 PM and interval length of 1 hour will be written at 1:00 PM.

**Daily** Represents a daily average of the performance of the resources. There will be one daily record written every day, per resource. Each daily record will be written at or shortly after midnight, according to the device clock. For example, a record with a timestamp of April 21, 12:00 AM and interval length of 24 hours will be written on April 22, 12:00 AM.

Note that hourly and daily records are usually retained in the database for longer periods of time, compared to By Sample records. As a result, if a generated report indicates that sample data is no longer available for a particular time range in the past, it is possible that the sample data has already been purged, in which case you can try generating the report using hourly or daily summation levels.

Be aware that when selecting a time range using the **Display historic performance data using relative time** radio button, it is necessary to select the length of the range in days. This can lead to unexpected results when also selecting the daily summation level. Remember that a daily record is not written until the day is complete, which means that at any given time during the day, the current day's data does not yet exist in the database. Therefore, selecting the daily summation level as well as a relative time range of **n days ago until now**, will usually result in **n-1** records being returned per resource, assuming that at least **n-1** days worth of data has been previously collected for those resources. Specifying a relative time range of **1 days ago until now** will result in an empty report.

## Available Columns, Included Columns

Use these items to specify which columns to include and which columns to exclude from the generated report.

### Available Columns

Displays the columns that you can include in the generated

report that are not already included. If a column appears in the **Included Columns** list box, it will not appear in the **Available Columns** list box.

### **Included Columns**

Displays the columns that will appear in the generated report. By default, all columns for a report are listed in this list box and will be displayed in the generated report. You can highlight the columns and use the up and down arrows to rearrange the order of the columns in the report. The order of the items in this list box determines the order in which the columns will appear in the generated report. The grayed-out column names shown at the top of the list box are the frozen columns, which cannot be selected, removed, or reordered for this type of report.

If a column appears in the **Available Columns** list box, it will not appear in the **Included Columns** list box.

To exclude a column from a report, click the name of the column in the **Included Columns** list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the left arrow button (<<). The selected columns are removed from the **Included Columns** list box, and will appear in the **Available Columns** list box.

To include a column in a report, click the name of the column in the **Available Columns** list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the right arrow button (>>). The selected columns will be removed from the **Available Columns** list box, and will appear in the **Included Columns** list box.

To rearrange the order of columns that will appear in the generated report, click the name of a column in the **Available Columns** list box. Click **Shift+click** and **Ctrl+click** to select multiple columns, if so desired. Then click the up arrow button (^) or the down arrow button (v) to move the selected columns up or down in the list, relative to the other columns.

5. Click the **Options** tab of the Batch Report job window, specify the following:
  - The machine on which the report file will be saved
  - The report format (for example, HTML)
  - Whether to run a script when the report process is complete
  - The format for the name of the batch report
6. In the **When to Run** tab of the Batch Report job window, specify when and how often to run the batch report.
7. In the **Alert** tab of the Batch Report job window, define the conditions that will trigger an alert and the method of notification.
8. Click **File > Save** and type a name for the batch report. The batch report will be saved using your user ID as a prefix. For example, if you are logged onto Tivoli Storage Productivity Center as "Smith," the name of the batch report will be `Smith.report_name`.

## **Editing batch reports**

Edit existing batch report definitions.

To edit a batch-report definition, follow this procedure:

1. Expand **IBM Tivoli Storage Productivity Center** → **Reporting** → **My Reports** → **Batch Reports**.
2. Click the report name, make your changes and then click **File** → **Save**.

## Viewing job information for a batch report

View job information for batch reports on the Job Management panel. Use this panel to check the status of a batch report schedule and view the runs, jobs, and log files associated with it.

To display a list of the runs and jobs for a batch report, complete the following steps:

1. Expand **IBM Tivoli Storage Productivity Center** → **Reporting** → **My Reports** → **Batch Reports**.
2. Click the name of a batch report. The content pane displays information about the batch report schedule.
3. Click **View Job History**. The Job Management panel is displayed and the batch report schedule is highlighted in the **Schedules** section.
4. View the runs and jobs for the highlighted batch report schedule in the **Jobs for Selected Schedule** section. See “Managing schedules, jobs, and runs” on page 295 for more information about how to use the Job Management panel.

## Displaying log files for a batch report

Display the log files batch reports.

To display the log files for a batch report schedule, complete these steps:

1. Expand **IBM Tivoli Storage Productivity Center** → **Reporting** → **My Reports** → **Batch Reports**.
2. Click the name of a batch report. The content pane displays information about the batch report schedule.
3. Click **View Job History**. The Job Management panel is displayed and the batch report schedule is highlighted in the **Schedules** section.
4. In the **Jobs for Selected Schedule**, expand the run whose log files you want to view. Click + or - to expand or collapse the job listing under a run.
5. Select the job whose log file you want to view and click **View Log File(s)**. See “Managing schedules, jobs, and runs” on page 295 for more information about how to use the Job Management panel.

## Displaying batch reports from the saved directory

Access the report file that is generated from a batch report schedule. This file contains the data for the report that you generated as a batch report.

After you run a batch report, batch report files are saved in the directory that you specified on the **Options** page of the Create Batch Report panel. Batch reports are automatically saved to the following default locations:

- Storage Resource agents:
  - Windows: C:\Program Files\IBM\TPC\agent\log\computername\reports
  - UNIX/Linux: /opt/IBM/TPC/agent/log/computername/reports
- Data agents:
  - Windows: C:\Program Files\IBM\TPC\ca\subagents\TPC\Data\log\computer\_name\reports\

- UNIX/Linux: /opt/IBM/TPC/ca/subagents/TPC/Data/log/computer\_name/reports/

## Viewing CSV batch reports with spreadsheet applications

View CSV-formatted batch reports in third-party spreadsheet applications.

If you are using a spreadsheet program (such as Excel) to view a CSV-type batch report file, that program might display data in the wrong columns either because the data was exported from Data Manager, the data contains commas, or the file was opened or formatted in the spreadsheet program by double-clicking it.

If data is displaying in the wrong columns, open the CSV file using the spreadsheet program's open command.

## Exporting and importing settings that you define

This topic describes how to export batch report definitions, constraints, groups, and profiles.

You can export settings that you define from one instance of IBM Tivoli Storage Productivity Center and import the settings into another instance. To export and import your definitions, complete the following steps:

1. On the system running the server component, change to the following directory:

|                   |                                             |
|-------------------|---------------------------------------------|
| On Linux and UNIX | /opt/IBM/TPC/Data/server/tools/             |
| On Windows        | C:\Program files\IBM\TPC\Data\server\tools\ |

2. To export the definitions from a command prompt, type the following command and press **Enter**:

```
export -e <file>
```

where <file> is the fully qualified name of the file where the definition information is written.

3. On the system to which you want to migrate the definitions, change to the **tools** directory.
4. To import the definitions, type the following command and press **Enter**:

```
export -i <file>
```

where <file> is the fully qualified name of the file that you created when exporting.

## User ID (saved) reports

Use the information in this section to learn how to generate, save, and view user ID reports. You can create a user ID reports based on a system report or any of the online reports available in Data Manager, Data Manager for Databases, Disk Manager, or Fabric Manager. Once saved, user ID reports are available in the **IBM Tivoli Storage Productivity Center > Reporting > My Reports > user ID name** node in the navigation tree.



## Creating a user-defined report

Create a custom report based on a system report or any of the online reports available in Data Manager, Data Manager for Databases, Disk Manager, or Fabric Manager.

To create a custom report, follow this procedure:

1. Specify new report criteria by doing one of the following:
  - Access a system report and change its format according to your organization's needs.
  - Create a new report using the reporting component in any of the IBM Tivoli Storage Productivity Center managers. For example, **Data Manager** → **Reporting** → **Usage** → **Access Load** → **Access Time** → **By Filesystem**.
2. Click **File** → **Save** and type a name for the report. The report is saved in the **Tivoli Storage Productivity Center** > **Reporting** > **My Reports** > *user name's* node, where *user name's* is your login name. For example, if you are logged in as "Smith", your custom reports appear under the **Smith's Reports** node.

## Displaying user-defined reports

Display a custom report saved under your user ID.

To display a list of custom reports, expand the **IBM Tivoli Storage Productivity Center** > **Reporting** > **My Reports** > *<your user ID>* node. Click the name of the report you want to view.

## Editing user-defined reports

Change the settings for an existing user ID report and save those changes.

To edit a custom report definition, follow this procedure:

1. Expand **My Reports** → *user name's Reports*.
2. Click the report name, make your changes and then click **File** → **Save**.

---

## Rollup reports

Use rollup reports to view the storage information collected by the master IBM Tivoli Storage Productivity Center server from its associated subordinate servers. Rollup reports enable you to have a network-wide perspective of storage usage in an environment where multiple Tivoli Storage Productivity Center servers are deployed and collecting storage metrics.

Before you can generate rollup reports, you must perform the following:

- Associate subordinate servers with the master server on the **Administrative Services** > **Data Sources** > **Tivoli Storage Productivity Center Servers** window.
- Run discoveries, probes, and scans from the subordinate servers to gather information about their locally-monitored storage entities.
- Run discoveries, probes, and scans from the master server to gather information about its locally-monitored storage entities. Note that information about the storage entities monitored by a master server are automatically shown in rollup reports.
- Run TPC Server probes from the master against the subordinate servers to collect their storage information for the rollup reports. Use the Tivoli Storage Productivity Center **Monitoring** > **TPC Server Probes** node to define and schedule probe jobs.

**Note:**

- If the same storage entity is managed by multiple subordinate servers, rollup reports reflect the storage information from the subordinate server that most recently probed that entity.
- After upgrading to Tivoli Storage Productivity Center v4.2 or later from a previous version, you must run probes against a master server's monitored storage assets to have information about those assets appear in **Tivoli Storage Productivity Center > Reporting > Rollup Reports**.

Tivoli Storage Productivity Center includes different rollup reporting categories that provide you with the flexibility to view data about your storage resources according to the needs of your environment. The following rollup reporting categories are available:

## Generating rollup reports

Learn how to generate and view rollup reports.

1. Expand the navigation tree to display the type of report you want to generate. For example, if you want to generate a database capacity report, expand **IBM Tivoli Storage Productivity Center > Reporting > Rollup Reports > Database Capacity**.
2. Highlight a By Clause to determine how you want to view a report. For example, to view rollup capacity information according to computer group, click **Tivoli Storage Productivity Center > Reporting > Rollup Reports > All DBMSs > By Computer Group**. The **Selection** page displays.
3. Use the **Selection** page to determine what columns appear within a report, what storage assets to report upon (Capacity and Database Capacity reports only), and how to filter that values that appear report columns.
4. (Capacity and Database Capacity reports only) Click **Selection...** to select the objects that you want to report upon from the Select Resources window. See "Selecting resources to appear on a report" on page 796 for information.
5. Click **Filter...** to further filter the objects that appear in a report. Filters enable you to apply general rules to the report based on the rows in that report. See "Filtering the resources that appear on a report" on page 799 for more information on filtering.
6. Use the **Available Columns** and **Included Columns** list boxes to determine what columns are displayed in a generated report.
7. Click **Generate Report**. A new tab will be added to the tab dialog representing the report that you generated.
8. Click the new tab to view the report. Reports are tabular in format and composed of rows and columns. You can scroll the report up and down and to the right and left to view the entire report. Use the **View** menu to hide/show the navigation tree to increase the viewable area of the report or drag the divider bar in the middle of the screen back and forth to reallocate the amount of space that is given each pane.

### Rollup reports - fragmentation

Fragmentation occurs when related storage entities are being monitored by different subordinate IBM Tivoli Storage Productivity Center servers.

The following example shows when fragmentation can occur using subordinate servers and rollup reports:

- **Example:** Host A is assigned a volume from storage subsystem 1. Host A is being monitored only by one Tivoli Storage Productivity Center subordinate

server (Server 1) and storage subsystem 1 is being monitored by another subordinate server (Server 2). Server 1 and Server 2 are subordinate servers to the master server (Server 3).

**Reason for fragmentation:** In this example, the host and subsystem are fragmented because they are not being monitored by the same Tivoli Storage Productivity Center server. This will cause the **Correlated Volume Space** value in the rollup reports for storage subsystems and volumes to be incorrect.

To prevent fragmentation when using master and subordinate server relationships to generate rollup reports, we recommend the following:

- Do not monitor a storage subsystem and the hosts to which its volumes are assigned with different subordinate servers that are associated with the same master server.
- Do not monitor a fabric and its related switches with different subordinate servers that are associated with the same master server.
- If an agent machine is using storage from a subsystem, then both the agent and subsystem should be monitored by the same subordinate server.
- Configure all the subordinate servers to discover all the storage subsystems within your environment. Note that if a subordinate server is monitoring a subsystem with an agent other than the agent that is using storage on that subsystem, the value for the **Correlated Volume Space** column in rollup reports will be incorrect. For example:
  - Host A has storage from Subsystem 1
  - Subordinate server B is Monitoring Host A and Subsystem 1
  - Subordinate server C is Monitoring Subsystem 1 and probed the subsystem more recently than server B.
  - The master server is monitoring subordinate servers B and server C.

The master server will report the storage subsystem data from server C. Because the master server does not know about Host A, the value for the **Correlated Volume Space** column will be incorrect.

## Rollup Reports - Asset

Use these reports to view detailed statistics about agents (IBM Tivoli Storage Productivity Center), computers, storage subsystems, disk and volume groups, disks, filesystems, logical volumes, volumes, and fabrics that are monitored by subordinate servers.

The information shown in these reports is gathered by the master server when it runs Tivoli Storage Productivity Center server probes against associated subordinate servers. If the same storage entity is managed by multiple subordinate servers, rollup reports reflect the storage information from the subordinate server that most recently probed that entity.

### Rollup Reports - Asset - Agents

Use this report to view information about Data agents, Fabric agents, and Storage Resource agents that are associated with subordinate servers in your environment.

You can generate and sort the data within this report using the following By clauses:

- By Agent: Shows agents sorted according to the name of the computer on which they are installed.

- **By OS Type:** Shows agents sorted according to the operating system under which they run. Click the magnifying glass icon next to an operating system type to view more detailed information about each of the agents running under that operating system.
- **By TPC Server:** Shows agents sorted according to the subordinate server that manages them. Click the magnifying glass icon next to subordinate server to view more detailed information about each of the agents associated with that server.

**Note:** The first column in each report remains displayed when scrolling through the other columns on the report.

## **By Agent**

### **Agent Name**

Name of the computer on which an agent is installed.

### **Data/SR Agent Version**

Tivoli Storage Productivity Center version number of a Data agent or Storage Resource agent (for example, 4.2.0).

### **Fabric Agent Version**

Tivoli Storage Productivity Center version number of a Fabric agent (for example, 4.1.1).

### **OS Type**

Operating system on the computer where an agent is installed.

### **Data/SR Agent Status**

Status of a Data agent or Storage Resource agent. Possible statuses include Up and Down. A down agent represents an agent that is still running but with which the server is not communicating.

### **Fabric Agent Status**

Status of a Fabric agent. Possible statuses include Up and Down. A down agent represents an agent that is still running but with which the server is not communicating.

### **Data/SR Agent Port**

Port number on which a Data agent or Storage Resource agent listens for requests. Port 9549 is registered for use by IBM.

### **Fabric Agent Port**

Port number on which a Fabric agent listens for requests. Port 9550 is registered for use by IBM.

### **Host Address**

Network name or IP address of an agent as seen by subordinate server.

### **Time Zone**

Time zone where the agent is installed (for example, CDT).

### **CPU Architecture**

Instruction set architecture on the computer where an agent is running.

## **By OS Type**

### **OS Type**

Operating system on the computer where an agent is installed.

**Number of Data Agents/SR Agents**

Total number of Data agents and Storage Resource agents that are running on the operating system.

**Number of Data Agents/SR Agents (Status: Up)**

Number of Data agents and Storage Resource agents that are running on the operating system that have a status of Up.

**Number of Data Agents/SR Agents (Status: Down)**

Number of Data agents and Storage Resource agents that are running on the operating system that have a status of Down. A down agent represents an agent that is still running but with which the server is not communicating.

**Number of Data Agents/SR Agents (Status: Defected)**

Number of Data agents and Storage Resource agents that are running on the operating system that have been reassigned to another server.

**Number of Data Agents/SR Agents (Status: Upgrading)**

Number of Data agents and Storage Resource agents that are running on the operating system and are currently in the process of upgrading to the same version as the subordinate server.

**Number of Data Agents/SR Agents (Status: Unreachable)**

Number of Data agents and Storage Resource agents that are running on the operating system that are unreachable. This status indicates that an agent that is unreachable and not responding to the server. This status might occur when the agent is not running but did not inform the server that it was shutting down, or when a network problem has occurred.

**Number of Data Agents/SR Agents (Status: Upgrade Required)**

Number of Data agents and Storage Resource agents that are running on the operating system that require an upgrade to the same version as the subordinate server.

**Number of Data Agents/SR Agents (Status: Deleted)**

Number of Data agents and Storage Resource agents that are running on the operating system that have a status of Deleted.

**Number of Fabric Agents**

Total number of Fabric agents that are running on the operating system.

**Number of Fabric Agents (Status: Up)**

Number of Fabric agents that are running on the operating system that have a status of Up.

**Number of Fabric Agents (Status: Down)**

Number of Fabric agents that are running on the operating system that have a status of Down. A down agent represents an agent that is still running but with which the server is not communicating.

**By TPC Server****TPC Server**

Name of the TPC server on which an agent is associated.

**Number of Data/SR Agents**

Total number of Data agents and Storage Resource agents that are associated with the TPC server.

**Number of Data Agents/SR Agents (Status: Up)**

Number of Data agents and Storage Resource agents that are associated with the TPC server that have a status of Up.

**Number of Data Agents/SR Agents (Status: Down)**

Number of Data agents and Storage Resource agents that are associated with the TPC server that have a status of Down. A down agent represents an agent that is still running but with which the server is not communicating.

**Number of Data Agents/SR Agents (Status: Defected)**

Number of Data agents and Storage Resource agents that are associated with the TPC server that have been reassigned to another server.

**Number of Data Agents/SR Agents (Status: Upgrading)**

Number of Data agents and Storage Resource agents that are associated with the TPC server and are currently in the process of upgrading to the same version as the subordinate server.

**Number of Data Agents/SR Agents (Status: Unreachable)**

Number of Data agents and Storage Resource agents that are associated with the TPC server that are unreachable. This status indicates that an agent that is unreachable and not responding to the server. This status might occur when the agent is not running but did not inform the server that it was shutting down, or when a network problem has occurred.

**Number of Data Agents/SR Agents (Status: Upgrade Required)**

Number of Data agents and Storage Resource agents that are associated with the TPC server that require an upgrade to the same version as the subordinate server.

**Number of Data Agents/SR Agents (Status: Deleted)**

Number of Data agents and Storage Resource agents that are associated with the TPC server that have a status of Deleted.

**Number of Fabric Agents**

Total number of Fabric agents associated with the TPC server.

**Number of Fabric Agents (Status: Up)**

Number of Fabric agents associated with the TPC server that have a status of Up.

**Number of Fabric Agents (Status: Down)**

Number of Fabric agents associated with the TPC server that have a status of Down.

**Rollup Reports - Asset - Computers**

Use this report to view information about the hosts and computers (including NetApp and NetWare) that are monitored by subordinate servers in your environment. This includes VMWare-related entities such as hypervisors and virtual machines.

**Note:** The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

You can generate and sort the data within this report using the following by clauses:

- By Computer: Shows computers monitored by subordinate servers. Sorted by computer name.

- By TPC Server: Shows subordinate servers that are monitoring computers. Sorted by subordinate server name. Click the magnifying glass icon next to subordinate server to view more detailed information about each of the computers associated with that server.

**Note:** The first column in each report remains displayed when scrolling through the other columns on the report.

## By Computer

**Note:** Two entries for the same host might appear in this report and the values for that host are counted twice in the **TOTAL** column. This scenario occurs under the following circumstances:

- You install two Storage Resource agents on the same host (using the long and short names of the host) that point to different Tivoli Storage Productivity Center servers. For example, you install an agent from serverA using the short name of a host (for example, hostC). You install another agent from serverB using the long name of the same host (for example, hostC.storage.chicago.com). You then set up the servers in a master and subordinate relationship. For example, serverA is the master and serverB is the subordinate. In this example, there are two entries in the report for the host: hostC and hostC.storage.chicago.com. The values for hostC are counted twice in the **TOTAL** column.
- A Storage Resource agent on a host (hostC) points to two different Tivoli Storage Productivity Center servers (serverA and serverB). The host has a dynamic IP address. You set up serverA as the master and serverB as the subordinate. Server serverA probes hostC, but then the IP address of hostC changes before serverB probes the host. In this example, the report might display two entries for hostC: one entry with the IP address that was probed by serverA; the other entry with the changed IP address that was probed by serverB. The values for hostC are counted twice in the **TOTAL** column.
- A Storage Resource agent on a host (hostC) points to two different Tivoli Storage Productivity Center servers (serverA and serverB). The host has a name that resolves to hostC.storage.chicago.com on serverA and hostC.srm.storage.chicago.com on serverB. You then set up the servers in a master and subordinate relationship. For example, serverA is the master and serverB is the subordinate. In this example, there are two entries in the report for the host: hostC.storage.chicago.com and hostC.srm.storage.chicago.com. The values for hostC are counted twice in the **TOTAL** column.

### Computer

Name of a computer monitored by a subordinate server.

### OS Type

Operating system running on a computer.

### OS Version

Version of the operating system running on a computer.

**RAM** Amount of RAM installed on a computer.

### Disk Space (Capacity)

Total storage capacity for a computer. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on a computer that has not been probed.

- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

**Note:** If the computer is an IBM Tivoli Storage SAN File Systems client, then this value will also include the capacity of any storage subsystem volumes visible to this computers that are owned by the SAN File System.

**Number of Disks**

Total number of disks for a computer. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value in this column:

- Computer disks on a computer that has not been probed by a Storage Resource agent or Data agent.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

**Processor Count**

Number of CPUs running on a computer.

**IP Address**

IP address of a computer.

**Serial Number**

Serial number of a computer.

**Model** Model number or name of a computer.

**Manufacturer**

Manufacturer of a computer (for example, IBM).

**Time Zone**

Time zone where the computer is located.

**Available File System Space (Filesystem Free Space)**

Total amount unique file system storage space that is not used or available to the computer. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

**Network Address**

Network name of a computer.

**Processor Type**

Type of processor running on a computer (for example, GenuineIntel: x86 Family 15).

**Processor Speed**

Speed at which a computer's processor is running (in MHz).

**CPU Arch**

Architecture of the processor (for example, IA32).

**Swap Space**

Amount of swap space (virtual memory) defined on a computer.



## By TPC Server

### TPC Server

Name of a subordinate server that is monitoring computers in your environment.

### # of Computers

Number of computers monitored by the subordinate server.

### Disk Space (Capacity)

Total storage capacity of the computers monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value for this column:

- Computer disks on a computer that has not been probed.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

**Note:** If a computer is an IBM Tivoli Storage SAN File Systems client, then this value will also include the capacity of any storage subsystem volumes visible to this computers that are owned by the SAN File System.

### Available File System Space (Filesystem Free Space)

Total amount unique file system storage space that is not used or available to the computers that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

### Number of Disks

Number of disks associated with the computers monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value in this column:

- Computer disks on a computer that has not been probed by the Storage Resource agent or Data agent.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

## Rollup Reports - Asset - Storage Subsystems

Use this report to view information about the storage subsystems that are monitored by subordinate servers in your environment.

**Note:** The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

You can generate and sort the data within this report using the following By clauses:

- By Storage Subsystem: Sorted according to storage systems name. The names of storage systems monitored by subordinate servers appear in the first column and remain displayed when scrolling through the other columns on the report.

- **By Disk Space:** Sorted according to the storage capacity of storage systems. The names of storage systems monitored by subordinate servers and their disk capacity appear in the first two columns and remain displayed when scrolling through the other columns on the report.
- **By TPC Server:** Shows storage systems aggregated at subordinate server level. The names of subordinate servers that are monitoring storage systems appear in the first column and remain displayed when scrolling through the other columns on the report. Click the magnifying glass icon next to subordinate server to view more detailed information about each of the storage systems monitored by that server.

**Note:**

- Storage systems that have been discovered but not probed are not displayed in these reports.
- This report includes information about storage system volumes. For IBM System Storage SAN Volume Controller systems, IBM Tivoli Storage Productivity Center uses the term "Volume" to represent "virtual disks" for those systems.

### **By Storage Subsystem, and By Disk Space**

**Name** Name or alias of a storage system that is monitored by a subordinate server or the master server.

**Disk Space (Capacity)**

Total raw (unformatted) disk capacity of a storage system. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Capacity of spare disks identified on IBM TotalStorage Enterprise Storage Server, IBM System Storage DS8000, and IBM System Storage DS6000 storage systems.
- Capacity of storage system disks that become missing after a storage system probe.

**Note:** For Hitachi Data Systems storage systems, this value is the capacity of all the volumes in the storage system.

**Available Disk Space (Unallocated Space)**

Total unformatted disk free space of the storage system. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this column:

- Storage system disks that become missing after a storage system probe.
- RAID overhead on storage systems that are identified by Tivoli Storage Productivity Center.
- External back-end storage for TagmaStore systems.
- Storage space that is consumed by volumes that reside on external disks when Hitachi Data Systems (HDS) storage systems are configured as storage virtualizers.

**Note:** For Hitachi Data Systems storage systems, internal disk information is not available. This field is set to 0.

**Number of Disks**

Total number of disks in the storage system. Tivoli Storage Productivity Center does not include disks that become missing after a successful

storage system probe in its calculation of the value for this column. For Hitachi Data Systems storage systems, internal disk information is not available. This field is to 1.

**Number of Volumes (LUNs)**

Number of storage volumes and virtual storage volumes discovered by a probe. Tivoli Storage Productivity Center does not include storage volumes or virtual storage volumes that become missing after a successful probe in its calculation of the value for this column.

**Volume Space (LUN Capacity)**

Total amount of unique storage system volume space or virtual storage volume space discovered on monitored storage system arrays. Tivoli Storage Productivity Center does not include storage volumes or virtual storage volumes that become missing after a successful probe in its calculation of the value for this column.

**Note:** For space-efficient volumes, this value is the amount of storage space requested for these volumes, not the actual allocated amount. This can result in discrepancies in the overall storage space reported for a storage system using space-efficient volumes. This also applies to other space calculations, such as the calculations for the storage system's Consumable Volume Space and FlashCopy Target Volume Space.

**Consumable Volume Space (Usable LUN Capacity)**

Total amount of unique storage system volume space or virtual storage volume space discovered on monitored storage system arrays. Tivoli Storage Productivity Center does not include the following storage entities in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a probe
- Storage system volumes or virtual storage volumes that are used as flash copy targets.
- Storage system volumes or virtual storage volumes that are identified as a Business Continuance Volume extended (BCVx) (EMC only).
- Storage system volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database (EMC only).

**FlashCopy Target Volume Space (Capacity)**

Total amount of unique storage system volume space and virtual storage volume space discovered on monitored storage system arrays and has been identified as flash copy target space. Tivoli Storage Productivity Center does not include the following storage entities in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a probe.
- Storage system volumes or virtual storage volumes that are not used as flash copy targets.

**Note:** TotalStorage Enterprise Storage Server systems must have at least the following microcode levels: `ess800mincodelevel = 2.4.3.56`, `essf20mincodelevel = 2.3.3.89`. Tivoli Storage Productivity Center does not report FlashCopy information if the TotalStorage Enterprise Storage Server systems do not meet this requirement.

**Correlated Volume Space (LUN Capacity Visible to Data Manager)**

Total amount of storage system volume space or virtual storage volume

space on systems where the systems disks are identified by Tivoli Storage Productivity Center as residing within the monitored storage systems in a group. Tivoli Storage Productivity Center does not include the following entities in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Disks that are assigned to a monitored storage system but are not identified by Tivoli Storage Productivity Center as a disk residing within the storage system. Reasons why these disks are not identified as residing within the storage system include:
  - Tivoli Storage Productivity Center does not support the identification of disks for some storage system types.
  - The connection of the disk to the storage system is lost and is no longer detectable during a probe.
- Storage system volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
- Storage system volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database.
- Storage system volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage system virtualizers. For example, IBM System Storage SAN Volume Controller.

#### **Volume Space Not Correlated (LUN Capacity Not Visible to Data Manager)**

Total amount of unique storage system volume space and virtual storage volume space discovered on monitored storage system arrays in a group that can be assigned or are assigned to systems within the network. Tivoli Storage Productivity Center does not include the following entities in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Disks that reside within a monitored storage system.
- Storage system volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage system virtualizers.

#### **Unavailable Disk Space (Overhead)**

Amount of storage system volume or virtual storage volume space that is dedicated to redundancy. This value is dependent on the storage system SMI-S provider returning a valid RAID value that can be used to determine the overhead. If the value is zero for a storage system, the overhead cannot be calculated. Tivoli Storage Productivity Center does not include the storage system volumes or virtual storage volumes that become missing after a probe in its calculation of the value for this column.

#### **Formatted Space (Capacity)**

Total amount of formatted storage space associated with the storage system that is used or can be used for system storage volumes or virtual storage volumes. This value is dependent on the storage system SMI-S provider returning a formatted value.

For some storage systems, this value is the total managed space which includes space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

For SAN Volume Controller, this value represents the total size of the managed disk groups created on the SAN Volume Controller. This value does not include managed disk space that is not defined to a managed disk group.

**Formatted Space with No Volumes (Formatted Space with No Volumes)**

Total amount of formatted storage space associated with the storage system that can be used for system storage volumes or virtual storage volumes. This value is dependent on the storage system SMI-S provider returning a formatted value.

For some systems, this value represents the remaining managed space that includes space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

For SAN Volume Controller, this value represents the space within the managed disk groups that does not include the size of the virtual disks created and the corresponding virtual disk extent overhead.

**Overall Unavailable Disk Space (Total Overhead)**

(TotalStorage Enterprise Storage Server, DS6000, DS8000, and SAN Volume Controller storage systems only.) Total amount of unformatted space within disk groups or array sites.

**Unformatted Disk Space (Undefined Disk Group or Array Site Capacity)**

This value does not apply to all storage systems. For TotalStorage Enterprise Storage Server, DS6000, DS8000 storage systems, this value is the total raw disk space of any unformatted disk groups or array sites. For IBM XIV Storage System, this value is the total available physical (hard) space that is left for pool creation. For other storage systems, this value is the total disk space that is not within a storage pool or N/A if the value is not applicable for the storage system.

**Cache** Amount of RAM or other cache for a storage system.

**Note:** This column is blank for SAN Volume Controllers.

**Manufacturer**

Manufacturer of a storage system.

**Model** Model number/name of a storage system.

**Serial Number**

Serial number of a storage system.

**Firmware Revision**

Microcode level of a storage system.

**Assigned Volume Space (Subsystem Assigned LUN Capacity)**

Total storage system volume space within the storage system that is mapped or assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a probe.
- Storage system volumes that are not mapped or are not assigned to host systems.

**Unassigned Volume Space (Subsystem Unassigned LUN Capacity)**

Total storage system volume space within a storage system that is not

mapped or not assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a probe.
- Storage system volumes that are mapped or assigned to host systems.

**Unassigned and Available Formatted Space (Subsystem Total Formatted Free Space)**

Total storage system space that is not mapped or not assigned to host systems and the total amount of formatted storage space associated with the storage system that can be used for system storage volumes.

This total amount of formatted storage space is dependent on the storage system SMI-S provider returning a formatted value. For some systems, this value represents the remaining managed space that might include space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes that become missing after a probe.
- Storage system volumes that are mapped or assigned to host systems.

**zOS Volume Space (LUN Capacity)**

Total storage system space that is reserved for use on mainframe storage such as used by zOS. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes that become missing after a storage system probe.
- Storage system volumes that are reserved for use on open system server or non-mainframe storage.

**Open System Volume Space (LUN Capacity)**

Total storage system space that is reserved for use on open system storage such as used by operating systems like Linux, AIX, HP/UX, Solaris, and Windows. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes that become missing after a storage system probe.
- Storage system volumes that are reserved for use on mainframe storage.

**Last Probed Timestamp**

Date and time when a storage system was last probed by the subordinate or master server.

**Physical Disk Space (Capacity)**

Total amount of physical disk space discovered on the monitored storage system arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system disks that become missing after a storage system probe.
- Back-end storage volumes on monitored storage system virtualizers (for example, SAN Volume Controller) that are attached from a storage system.

**Backend Volume Space (LUN Capacity)**

Total amount of storage system volume space or virtual storage space that

Tivoli Storage Productivity Center identifies as the back-end disk space on a monitored storage system virtualizer (for example, SAN Volume Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes or virtual storage volumes that are back-end disks on a storage system virtualizer that is not monitored.

## **By TPC Server**

### **TPC Server**

Name of a subordinate server that is monitoring storage systems in your environment.

### **Number of Subsystems (Total)**

Number of storage systems monitored by the subordinate server.

### **Disk Space (Total Disk Capacity)**

Total raw (unformatted) disk capacity of the storage systems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Capacity of spare disks identified on TotalStorage Enterprise Storage Server, DS8000, and DS6000 storage systems.
- Capacity of storage system disks that become missing after a storage system probe.

**Note:** For Hitachi Data Systems storage systems, this value is the capacity of all the Volumes in the storage system.

### **Available Disk Space (Total Disk Unallocated Space)**

Total unformatted disk free space of the storage systems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following storage space in its calculation of the value for this column:

- Storage system disks that become missing after a storage system probe.
- RAID overhead on storage systems that are identified by Tivoli Storage Productivity Center.
- External back-end storage for TagmaStore systems.
- Storage space that is consumed by volumes that reside on external disks when Hitachi Data Systems (HDS) storage systems are configured as storage virtualizers.

**Note:** For Hitachi Data Systems storage systems, internal disk information is not available. This field is set to 0.

### **Number of Disks (Total)**

Total number of disks in the storage systems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include disks that become missing after a successful storage system probe in its calculation of the value for this column. For Hitachi Data Systems storage systems, internal disk information is not available. This field is set to 1.

### **Number of Volumes (Total Number of LUNs)**

Total number of storage volumes and virtual storage volumes discovered by the subordinate server. Tivoli Storage Productivity Center does not

include storage volumes or virtual storage volumes that become missing after a successful probe in its calculation of the value for this column.

#### **Volume Space (Total LUN Capacity)**

Total amount of unique storage system volume space or virtual storage volume space discovered on storage system arrays that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include storage volumes or virtual storage volumes that become missing after a successful probe in its calculation of the value for this column.

**Note:** For space-efficient volumes, this value is the amount of storage space requested for these volumes, not the actual allocated amount. This can result in discrepancies in the overall storage space reported for a storage system using space-efficient volumes. This also applies to other space calculations, such as the calculations for the storage system's Consumable Volume Space and FlashCopy Target Volume Space.

#### **Consumable Volume Space (Total Usable LUN Capacity)**

Total amount of unique storage system volume space or virtual storage volume space discovered on storage system arrays that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following storage entities in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a probe
- Storage system volumes or virtual storage volumes that are used as flash copy targets.
- Storage system volumes or virtual storage volumes that are identified as a Business Continuance Volume extended (BCVx) (EMC only).
- Storage system volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database (EMC only).

#### **FlashCopy Target Volume Space (Total FlashCopy Target Capacity)**

Total amount of unique storage system volume space and virtual storage volume space discovered on storage system arrays that are monitored by the subordinate server and have been identified as flash copy target space. Tivoli Storage Productivity Center does not include the following storage entities in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a probe.
- Storage system volumes or virtual storage volumes that are not used as flash copy targets.

**Note:** TotalStorage Enterprise Storage Server systems must have at least the following microcode levels: `ess800mincodelevel = 2.4.3.56`, `essf20mincodelevel = 2.3.3.89`. Tivoli Storage Productivity Center does not report FlashCopy information if the TotalStorage Enterprise Storage Server systems do not meet this requirement.

#### **Unavailable Disk Space (Overhead)**

Amount of storage system volume or virtual storage volume space that is dedicated to redundancy and is monitored by the subordinate server. This value is dependent on the storage system SMI-S provider returning a valid RAID value that can be used to determine the overhead. If the value is zero for a storage system, the overhead cannot be calculated. Tivoli Storage



Productivity Center does not include the storage system volumes or virtual storage volumes that become missing after a probe in its calculation of the value for this column.

#### **Correlated Volume Space (Total LUN Capacity Visible to Data Manager)**

Total amount of storage system volume space or virtual storage volume space on systems where the systems disks are identified by Tivoli Storage Productivity Center as residing within the storage systems that are monitored by the subordinate server. The value in this column reflects the systems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following entities in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Disks that are assigned to a monitored storage system but are not identified by Tivoli Storage Productivity Center as a disk residing within the storage system. Reasons why these disks are not identified as residing within the storage system include:
  - Tivoli Storage Productivity Center does not support the identification of disks for some storage system types.
  - The connection of the disk to the storage system is lost and is no longer detectable during a probe.
- Storage system volumes or virtual storage volumes that are identified as Business Continuance Volume extended (BCVx).
- Storage system volumes or virtual storage volumes used as a Volume Configuration Management (VCM) database.
- Storage system volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage system virtualizers. For example, SAN Volume Controller.

#### **Volume Space Not Correlated (Total LUN Capacity Not Visible to Data Manager)**

Total amount of unique storage system volume space and virtual storage volume space discovered on monitored storage system arrays that can be assigned or are assigned to servers within the network. The value in this column reflects the storage systems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following entities in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Disks that reside within a monitored storage system.
- Storage system volumes or virtual storage volumes that Tivoli Storage Productivity Center identifies as back-end storage volumes on monitored storage system virtualizers.

#### **Overall Unavailable Disk Space (Total Formatted Capacity)**

Total amount of formatted storage space that is used or can be used for system storage volumes or virtual storage volumes. The value in this column reflects the storage systems that are monitored by the subordinate server. This value is dependent on the storage system SMI-S provider returning a formatted value.

For some storage systems, this value is the total managed space which includes space that cannot be used for storage volumes or virtual storage

volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

**Note:** This value does not apply to SAN Volume Controller.

**Formatted Space (Total Formatted Space with No Volumes)**

Total amount of formatted storage space that can be used for system storage volumes or virtual storage volumes. The value in this column reflects the storage systems that are monitored by the subordinate server. This value is dependent on the storage system SMI-S provider returning a formatted value.

For some systems, this value represents the remaining managed space that includes space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

For SAN Volume Controller, this value represents the total size of the managed disk groups created on the SAN Volume Controller. This value does not include managed disk space that is not defined to a managed disk group.

**Formatted Space with No Volumes (Total Subsystem Total Formatted Free Space)**

Total storage system space that is not mapped or not assigned to host systems and the total amount of formatted storage space associated with the storage system that can be used for system storage volumes. The value in this column reflects the storage systems that are monitored by the subordinate server.

This total amount of formatted storage space is dependent on the storage system SMI-S provider returning a formatted value. For some systems, this value represents the remaining managed space that might include space that cannot be used for storage volumes or virtual storage volumes. This can occur due to a number of reasons, such as space not being formatted until storage volumes are created.

For SAN Volume Controller, this value represents the space within the managed disk groups that does not include the size of the virtual disks created and the corresponding virtual disk extent overhead.

Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes that become missing after a probe.
- Storage system volumes that are mapped or assigned to host systems.

**Capacity Overhead (Total)**

(TotalStorage Enterprise Storage Server and SAN Volume Controller only.)  
Sum of all the space in a disk group minus the formatted capacity of the storage systems monitored by a subordinate server.

**Unformatted Disk Space (Total Undefined Disk Group or Array Site Capacity)**

This value does not apply to all storage systems. For TotalStorage Enterprise Storage Server, DS6000, DS8000 storage systems, this value is the total raw disk space of any unformatted disk groups or array sites that are monitored by the subordinate server. For the XIV system, this value is the total available physical (hard) space that is left for pool creation. For other storage systems, this value is the total disk space that is not within a storage pool or **N/A** if the value is not applicable for the storage system.

**Cache (Total)**

Amount of RAM or other cache for the storage systems that are monitored by the subordinate server.

**Assigned Volume Space**

Total storage system volume space within the storage system that is mapped or assigned to host systems. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes that are not mapped or are not assigned to host systems.

**Unassigned Volume Space (Total Subsystem Assigned LUN Capacity)**

Total storage system volume space within storage systems that is mapped or assigned to host systems. The value in this column reflects the storage systems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a probe.
- Storage system volumes that are not mapped or are not assigned to host systems.

**Unassigned and Available Formatted Disk Space (Total Subsystem Unassigned LUN Capacity)**

Total storage system volume space within storage systems that is not mapped or not assigned to host systems. The value in this column reflects the storage systems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a probe.
- Storage system volumes that are mapped or assigned to host systems.

**zOS Volume Space (Total zOS LUN Capacity)**

Total storage system space that is reserved for use on mainframe storage such as used by zOS. The value in this column reflects the storage systems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes that become missing after a storage system probe.
- Storage system volumes that are reserved for use on open system server or non-mainframe storage.

**Open System Volume Space (Total Open System LUN Capacity)**

Total storage system space that is reserved for use on open system storage such as used by operating systems like Linux, AIX, HP/UX, Solaris, and Windows. The value in this column reflects the storage systems that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes that become missing after a storage system probe.

- Storage system volumes that are reserved for use on mainframe storage.

### **Physical Disk Space (Capacity)**

Total amount of physical disk space discovered on the storage system arrays that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system disks that become missing after a storage system probe.
- Back-end storage volumes on monitored storage system virtualizers (for example, SAN Volume Controller) that are attached from a storage system.

### **Backend Volume Space (LUN Capacity)**

Total amount of storage system volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on a monitored storage system virtualizer (for example, SAN Volume Controller). The value in this column reflects the volumes that are monitored by the subordinate server. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes or virtual storage volumes that are back-end disks on a storage system virtualizer that is not monitored.

## **Rollup Reports - Asset - Disk/Volume Groups**

Use this report to view information about the disk and volume groups that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- **By Computer:** Shows disk and volume groups aggregated at the computer level and sorted by computer name.
- **By Disk/Volume Group Space:** Shows disk and volume groups aggregated at the computer level and sorted by storage capacity.
- **By Available Disk/Volume Group Space:** Shows disk and volume groups aggregated at the computer level and sorted by storage free space.
- **By TPC Server:** Shows disk and volume groups aggregated at the subordinate server level and sorted by subordinate server name. Click the magnifying glass icon next to the name of a subordinate server to view information about the computers it monitors.

**Note:** The first column in each report remains displayed when scrolling through the other columns on the report.

### **By Computer**

**Note:** The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

#### **Computer Name**

Name of a computer with disk and volume groups that is monitored by a subordinate server.

**Number of Disk/Volume Groups (Volume Groups)**

Number of volume groups on a computer.

**Disk/Volume Group Space (Total Capacity)**

Total storage capacity of the disks and volume groups on a computer.

**Available Disk/Volume Group Space (Total Free Space)**

Amount of storage space available on the disk and volume groups that are monitored by a subordinate server.

**OS Type**

Operating system running on a computer.

**Number of Disks**

Number of disks on a computer.

**Last Probe Time**

Date and time when a computer was last probed.

**By Disk/Volume Group Space (Freespace)****Computer Name**

Name of a computer with disk and volume groups that are monitored by a subordinate server.

**Disk/Volume Group Space (Total Free Space)**

Amount of storage space available on the disks and volume groups associated with a computer.

**Number of Disk/Volume Groups (Number of Volume Groups)**

Number of volume groups on a computer.

**Available Disk/Volume Group Space (Total Capacity)**

Total storage capacity of the disks and volume groups on a computer.

**OS Type**

Operating system running on a computer.

**Number of Disks**

Number of disks on a computer.

**Last Probe Time**

Date and time when a computer was last probed.

**By Available Disk/Volume Group Space (Capacity)****Computer Name**

Name of a computer with disk and volume groups that is monitored by a subordinate server.

**Available Disk/Volume Group Space (Total Capacity)**

Total storage capacity of the disks and volume groups on a computer.

**Number of Disk/Volume Groups (Volume Groups)**

Number of volume groups on a computer.

**Disk/Volume Group Space (Total Free Space)**

Amount of storage space available on the disks and volume groups associated with a computer.

**OS Type**

Operating system running on a computer.

**Number of Disks**

Number of disks on a computer.

**Last Probe Time**

Date and time when a computer was last probed.

**By TPC Server****TPC Server**

Name of a subordinate server that is monitoring computers with disk and volume groups.

**Number of Disk/Volume Groups**

Number of disk and volume groups that are monitored by a subordinate server.

**Disk/Volume Group Space (Total Free Space)**

Amount of storage space available on the disks and volume groups that are monitored by a subordinate server.

**Available Disk/Volume Group Space (Total Capacity)**

Total storage capacity of the disks and volume groups that are monitored by a subordinate server.

**Rollup Reports - Asset - Disks**

Use this report to view information about the disks that are monitored by subordinate servers in your environment.

You can generate and sort the data within Disk reports using the following by clauses:

- **By Available Disk Space (Computers):** Shows disks aggregated at the computer level and sorted by unallocated space. The names of computers and their system type appear in the first two columns and remain displayed when scrolling through the other columns on the report.
- **By Available Disk Space (Subsystems):** Shows disks aggregated at the storage subsystem level and sorted by unallocated space. The names of storage subsystems and their system type appear in the first two columns and remain displayed when scrolling through the other columns on the report.
- **By TPC Server (computers):** Shows computer disks aggregated at the subordinate server level. The names of subordinate servers appear in the first column and remain displayed when scrolling through the other columns on the report. Click the magnifying glass icon next to a subordinate server name to view more information about the computers monitored by that subordinate server.
- **By TPC Server (subsystems):** Shows storage subsystem disks aggregated at the subordinate server level. The names of subordinate servers appear in the first column and remain displayed when scrolling through the other columns on the report. Click the magnifying glass icon next to a subordinate server name to view more information about the storage subsystems monitored by that subordinate server.

**Note:** The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

**By Available Disk Space (Computers)****TOTAL**

Shows the system-wide totals for each of the columns in the report.

**System Name**

Name of a computer with disks that are monitored by a subordinate server.

**OS Type**

Operating system running on a computer.

**Available Disk Space (Total Unallocated Space)**

Amount of storage space on disks that is not allocated to any logical volume. Note: the value that appears in the TOTAL row for this column includes the unallocated capacity of all disks within the computers monitored by subordinate servers.

**Disk Space (Total Capacity)**

Amount of storage space available on the disks associated with a computer.

**Number of Disks**

Number of disks on a computer.

**By Available Disk Space (Subsystems)****System Name**

Name of a storage subsystem with disks that are monitored by a subordinate server.

**Subsystem Type**

Operating system running on a storage subsystem.

**Available Disk Space (Total Unallocated Space)**

Amount of space on the disk that is not allocated to any volumes for disks within storage subsystems). Note: the value that appears in the TOTAL row for this column includes the unallocated capacity of all disks within the storage subsystem, including the capacity of spare disks.

**Disk Space (Total Capacity)**

Amount of storage space available on the disks associated with a storage subsystem.

**Note:** For space-efficient volumes, this value is the amount of storage space requested for these volumes, not the actual allocated amount. This can result in discrepancies in the overall storage space reported for a storage system using space-efficient volumes. This also applies to other space calculations, such as the calculations for the storage system's Consumable Volume Space and FlashCopy Target Volume Space.

**Number of Disks**

Total number of disks in a storage subsystem.

**Physical Disk Space**

Total amount of physical disk space discovered on the monitored storage subsystem arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem disks that become missing after a storage subsystem probe.
- Back-end storage volumes on monitored storage subsystem virtualizers (for example, IBM SAN Volume Controller) that are attached from a storage subsystem.

## By TPC Server (Computers)

### TOTAL

Shows the system-wide totals for each of the columns in the report.

### TPC Server

Name of a subordinate server that is monitoring disks on computers.

### Available Disk Space (Total Unallocated Space)

Amount of storage space on computer disks that is not allocated to any logical volume and are monitored by a subordinate server.

### Disk Space (Total Capacity)

Amount of storage space available on all the computer disks that are monitored by a subordinate server.

**Note:** For space-efficient volumes, this value is the amount of storage space requested for these volumes, not the actual allocated amount. This can result in discrepancies in the overall storage space reported for a storage system using space-efficient volumes. This also applies to other space calculations, such as the calculations for the storage system's Consumable Volume Space and FlashCopy Target Volume Space.

### Number of Disks

Number of computer disks monitored by a subordinate server.

## By TPC Server (Subsystems)

### TPC Server

Name of a subordinate server that is monitoring disks on storage subsystems.

### Available Disk Space (Total Unallocated Space)

Amount of storage space on storage subsystem disks that is not allocated to any volume and are monitored by a subordinate server.

### Disk Space (Total Capacity)

Amount of storage space available on all the storage subsystem disks that are monitored by a subordinate server.

### Number of Disks

Number of storage subsystem disks monitored by a subordinate server.

### Physical Disk Space

Total amount of physical disk space discovered on the monitored storage subsystem arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem disks that become missing after a storage subsystem probe.
- Back-end storage volumes on monitored storage subsystem virtualizers (for example, IBM SAN Volume Controller) that are attached from a storage subsystem.

## Rollup Reports - Asset - File Systems or Logical Volumes

Use this report to view information about the filesystems or logical volumes that are monitored by subordinate servers in your environment.

You can generate and sort the data within File Systems or Logical Volumes reports using the following by clauses:



- By Available File System or Logical Volume Space: Shows filesystems aggregated at the computer level and sorted by free space and computer name.
- By TPC Server: Shows filesystems aggregated at the subordinate server level and sorted by subordinate server name. Click the magnifying glass icon next to a subordinate server name to view more information about the filesystems or logical volumes monitored by that subordinate server.

**Note:** The first column in each report remains displayed when scrolling through the other columns on the report.

**Note:** The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

## **By Available File System or Logical Volume Space**

### **Computer Name**

Name of a computer with filesystems or logical volumes that are monitored by a subordinate server.

### **Available File System/Logical Volume Space (Free Space)**

Amount of storage space available in the filesystems or logical volumes on a computer.

### **Consumed File System/Logical Volume Space (Used Space)**

Amount of storage space consumed by the filesystems or logical volumes on a computer.

### **File System/Logical Volume Space (Capacity)**

Total storage capacity of the filesystems or logical volumes on a computer.

### **OS Type**

Operating system running on a computer.

### **Number of File Systems/Logical Volumes (Number of FSs/LVs)**

Number of filesystems or logical volumes on a computer.

## **By TPC Server**

### **TPC Server**

Name of a subordinate server that is monitoring the filesystems and logical volumes in your environment.

### **Available File System/Logical Volume Space (Total Free Space)**

Amount of storage space available in all the filesystems or logical volumes on the computers monitored by the subordinate server.

### **Consumed File System/Logical Volume Space (Total Used Space)**

Amount of storage space consumed by all the filesystems or logical volumes on the computers monitored by the subordinate server.

### **File System/Logical Volume Space (Total Capacity)**

Total storage capacity of all the filesystems or logical volumes on the computers monitored by the subordinate server.

### **Number of File Systems/Logical Volumes (Number of FSs/LVs)**

Number of filesystems or logical volumes on the computers monitored by the subordinate server.

## Rollup Reports - Asset - Storage Subsystem Volumes

Use this report to view information about the Storage Subsystem Volumes that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Storage Subsystem: Shows Storage Subsystem Volume capacity and count aggregated at the storage subsystem level.
- By Storage Subsystem Volume Space: Shows Storage Subsystem Volume capacity and count aggregated at the subsystem level. Sorted according to the storage capacity of Storage Subsystem Volumes.
- By TPC Server: Shows Storage Subsystem Volume capacity and count aggregated at the subordinate server level. Sorted by subordinate server name. Click the magnifying glass icon next to a subordinate server to view information about the storage subsystems it monitors.

**Note:** The first column in each report remains displayed when scrolling through the other columns on the report.

**Note:** The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

### By Storage Subsystem

#### Storage Subsystem Name

Name of a storage subsystem containing Storage Subsystem Volumes that is monitored by a subordinate server.

#### Storage Subsystem Type

Type of storage subsystem, such as the IBM SAN Volume Controller.

#### Volume Space (LUN Capacity)

Total storage capacity of the Storage Subsystem Volumes in a storage subsystem.

#### Consumable Volume Space (Usable LUN Capacity)

Storage capacity of Storage Subsystem Volumes in the created disk array (not including BCVs, VCM database, and Replication targets).

#### Correlated Volume Space (LUN Capacity Visible to Data Manager)

Capacity of all Storage Subsystem Volumes visible to a Storage Resource agent or Data agent associated with a storage subsystem.

#### Formatted Space (Formatted Capacity)

(IBM TotalStorage Enterprise Storage Server, IBM FASTT, HP StorageWorks, and Hitachi Data Systems only) Total amount of formatted storage space in the storage pools associated with a storage subsystem. For example, a RAID volume space has more capacity than a stripe-mirror volume space.

For SAN Volume Controller, this value represents the total size of the managed disk groups created on the SAN Volume Controller. This value does not include managed disk space that is not defined to a managed disk group.

#### Assigned Volume Space (Subsystem Assigned LUN Capacity)

Total storage capacity of all Storage Subsystem Volumes that the storage subsystem knows are assigned to a host computer.

**zOS Volume Space (zOS LUN Capacity)**

Total storage capacity for all the Storage Subsystem Volumes allocated to zOS. Data for this column is calculated for TotalStorage Enterprise Storage Server only. The value for this column is set to N/A for other storage subsystems.

**Number of Volumes (Number of LUNs)**

Number of Storage Subsystem Volumes associated with a storage subsystem.

**Backend Volume Space**

Total amount of storage subsystem volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on a monitored storage subsystem virtualizer (for example, SAN Volume Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are back-end disks on a storage subsystem virtualizer that is not monitored.

**By Storage Subsystem Volume Space****Volume Space (LUN Capacity)**

Total storage capacity of the Storage Subsystem Volumes in a storage subsystem.

**Consumable Volume Space (Usable LUN Capacity)**

Storage capacity of Storage Subsystem Volumes in the created disk array (not including BCVs, VCM database, and Replication targets).

**Correlated Volume Space (LUN Capacity Visible to Data Manager)**

Capacity of all Storage Subsystem Volumes visible to a Storage Resource agent or Data agent associated with a storage subsystem.

**Formatted Space (Formatted Capacity)**

(IBM TotalStorage Enterprise Storage Server, IBM FAStT, HP StorageWorks, and Hitachi Data Systems only) Total amount of formatted storage space in the storage pools associated with a storage subsystem. For example, a RAID volume space has more capacity than a stripe-mirror volume space.

For SAN Volume Controller, this value represents the total size of the managed disk groups created on the SAN Volume Controller. This value does not include managed disk space that is not defined to a managed disk group.

**Assigned Volume Space (Subsystem Assigned LUN Capacity)**

Total storage capacity of all Storage Subsystem Volumes that the storage subsystem knows are assigned to a host computer.

**zOS Volume Space (zOS LUN Capacity)**

Total storage capacity for all the storage subsystem volumes allocated to zOS. Data for this column is calculated for TotalStorage Enterprise Storage Server and SAN Volume Controller and SAN Volume Controller only. The value for this column is set to N/A for other storage subsystems.

**Number of Volumes (Number of LUNs)**

Number of storage subsystem volumes associated with a storage subsystem.

**Storage Subsystem Name**

Name of a storage subsystem containing storage subsystem volumes that is monitored by a subordinate server.

**Storage Subsystem Type**

Type of storage subsystem, such as the IBM SAN Volume Controller.

**Backend Volume Space**

Total amount of storage subsystem volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the backend disk space on a monitored storage subsystem virtualizer (for example, SAN Volume Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are backend disks on a storage subsystem virtualizer that is not monitored.

**By TPC Server****TPC Server**

Name of a subordinate server that is monitoring storage subsystems.

**Storage Subsystem Count**

Number of storage subsystems that are monitored by a subordinate server.

**Volume Space (LUN Capacity)**

Total storage capacity of the Storage Subsystem Volumes that are monitored by a subordinate server.

**Consumable Volume Space (Usable LUN Capacity)**

Storage capacity of Storage Subsystem Volumes in the created disk array (not including BCVs, VCM database, and Replication targets) that are monitored by a subordinate server.

**Correlated Volume Space (LUN Capacity Visible to Data Manager)**

Capacity of all Storage Subsystem Volumes visible to a Storage Resource agent or Data agent associated with a subordinate server.

**Formatted Space (Formatted Capacity)**

(IBM TotalStorage Enterprise Storage Server, IBM FAStT, HP StorageWorks, and Hitachi Data Systems only) Total amount of formatted storage space in the storage pools associated with the storage subsystems that are monitored by a subordinate server. For example, a RAID volume space has more capacity than a stripe-mirror volume space.

For SAN Volume Controller, this value represents the total size of the managed disk groups created on the SAN Volume Controller. This value does not include managed disk space that is not defined to a managed disk group.

**Assigned Volume Space (Subsystem Assigned LUN Capacity)**

Total storage capacity of all Storage Subsystem Volumes that the monitored storage subsystems know are assigned to a host computer.

**zOS Volume Space (zOS LUN Capacity)**

Total storage capacity for all the Storage Subsystem Volumes allocated to zOS that are monitored by a subordinate server. Data for this column is calculated for TotalStorage Enterprise Storage Server and SAN Volume Controller only. The value for this column is set to N/A for other storage subsystems.

**Number of Volumes (Number of LUNs)**

Number of Storage Subsystem Volumes associated with the storage subsystems monitored by a subordinate server.

**Backend Volume Space**

Total amount of storage subsystem volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the backend disk space on a monitored storage subsystem virtualizer (for example, SAN Volume Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage subsystem volumes or virtual storage volumes that become missing after a storage subsystem probe.
- Storage subsystem volumes or virtual storage volumes that are backend disks on a storage subsystem virtualizer that is not monitored.

**Rollup Reports - Asset - Fabrics**

Use this report to view information about the fabrics that are monitored by subordinate servers in your environment.

You can view the following reports about fabrics:

- SAN Assets (ALL): Use this report to view information about all the SAN assets for each subordinate server.
- SAN Assets (Connected Dev): Use this report to view information about all the connected devices for each subordinate server.
- SAN Assets (Switches): Use this report to view information about switches in your SAN for each subordinate server.

**SAN Assets (ALL)****TPC Server**

Name of a subordinate server that is monitoring fabrics.

**Fabric ID**

The ID of the fabric that the asset is in.

**Type** The type of asset discovered. For example: switch, host, host bus adapter, storage subsystem.

**Label** The label of the asset. For example: the host name, switch name, storage subsystem name, node name of the port.

**Identifier**

The unique identifier of the switch. For example, the switch WWN.

**Status** The status of the asset: Normal, Warning, or Critical.

**SAN Assets (Connected Dev)****TPC Server**

Name of a subordinate server that is monitoring a connected device.

**Fabric ID**

The ID of the fabric that the asset is in.

**Type** The type of asset discovered. For example: switch, host, host bus adapter, storage subsystem.

**Label** The label of the asset. For example: the host name, switch name, storage subsystem name, node name of the port.

**Identifier**

The unique identifier of the switch. For example, the switch WWN.

**Status** The status of the asset: Normal, Warning, or Critical.

**SAN Assets (Switches)****TPC Server**

Name of a subordinate server that is monitoring a switch.

**Fabric ID**

The ID of the fabric that a switch is in.

**Type** The type of switch discovered.

**Switch Type**

The switch type, such as physical.

**Label** The label of a switch.

**Status** The status of a switch: Normal, Warning, or Critical.

**IP Address**

The IP address of a switch.

**Vendor ID**

The manufacturer of the switch.

**Model** The model name and number of a switch.

**Version**

The version information of the switch.

**Serial Number**

The serial number of the switch.

**Object ID**

The object ID of the switch.

**Management ID**

The management ID of the switch.

**Management Address**

The URL of the switch element manager.

**Domain**

The name of the domain that the switch is in.

**Contact**

The contact information for the switch.

**WWN** The world-wide name (WWN) of the switch.

**Location**

The physical location of the switch.

**Physical Switch WWN**

The WWN for the parent or physical switch of a virtual switch. If this is not a virtual switch then this column is empty.

**Rollup Reports - Asset - Hypervisors**

Use this report to view information about the hypervisors that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clause:

- By Hypervisor: Sorted according to hypervisor name. The names of hypervisors monitored by subordinate servers appear in the first column and remain displayed when scrolling through the other columns on the report.

## **By Hypervisors**

### **Hypervisor Name**

Name of a hypervisor monitored by a subordinate server.

### **OS Version**

Version of the operating system running on a computer.

**RAM** The amount of RAM installed on a computer.

### **Disk Space**

Total raw (unformatted) disk capacity of a storage subsystem. Tivoli Storage Productivity Center does not include the following capacity information in its calculation of the value for this column:

- Capacity of spare disks identified on IBM TotalStorage Enterprise Storage Server, IBM DS8000, and IBM DS6000 storage subsystems.
- Capacity of storage subsystem disks that become missing after a storage subsystem probe.

**Note:** For Hitachi Data Systems storage subsystems, this is the capacity of all the Volumes in the storage subsystem.

### **Number of Disks**

Total number of disks for a computer. Tivoli Storage Productivity Center does not include the following disks in its calculation of the value in this column:

- Computer disks on a computer that has not been probed by a Storage Resource agent or Data agent.
- Computer disks that become missing after a probe. Disks can become undetectable for a number of reasons, including a physical removal of the disks from a computer.

### **Processor Count**

Number of CPUs running on a computer.

### **IP Address**

IP address of a computer.

### **Serial Number**

Serial number of a computer.

**Model** Model number or name of a computer.

### **Manufacturer**

Manufacturer of a computer (for example, IBM).

### **Time Zone**

Time zone where the computer is located.

### **Available File System Space**

Total amount unique file system storage space that is not used or available to the computer. Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.

- File systems that are not mounted.

**OS Type**

Operating system running on a computer.

**Network**

Network name of a computer.

**Processor Type**

Type of processor running on a computer (for example, GenuineIntel: x86 Family 15).

**Processor Speed**

Speed at which a computer's processor is running (in MHz).

**CPU Architecture**

Architecture of the processor (for example, IA32).

**Swap Space**

Amount of swap space (virtual memory) defined on a computer.

## Rollup Reports - Database Asset

Use these reports to view detailed statistics that have been collected by subordinate servers about the RDBMSs in your environment, including Oracle, SQL Server, Sybase, and UDB/DB2.

The information shown in these reports is gathered by the master server when it runs Tivoli Storage Productivity Center server probes against associated subordinate servers. If the same storage entity is managed by multiple subordinate servers, rollup reports reflect the storage information from the subordinate server that most recently probed that entity.

### Rollup Reports - Database Asset - All DBMSs

Use this report to view information about the RDBMSs (including Oracle, Sybase, SQL Server, and UDB/DB2) that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Instance: Shows databases for each computer sorted by computer name.
- By Version: Shows databases for each computer sorted by RDBMS version.
- By DBMS Type: Shows databases for each computer sorted by RDBMS type.
- By TPC Server: Shows databases aggregated at the subordinate server level and sorted by subordinate server name.

**Note:** The first column in each report remains displayed when scrolling through the other columns on the report.

**Note:** The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

#### By Instance

**Computer Name**

Name of the computer on which an Instance resides.

**Instance**

Name of an instance.



**RDBMS Type**

RDBMS to which an instance belongs (e.g., Oracle, SQL Server, Sybase, UDB, etc.).

**RDBMS Version**

Version of the RDBMS under which an instance is running.

**Port** Port on which an instance listens for requests.

**Home Directory**

Path of the home directory for an instance.

**Logonid**

ID used by Tivoli Storage Productivity Center to log in to an instance.

**Probe Status**

Status of the last probe run against an instance.

**Last Start Time**

Date and time when an instance was last started.

**Last Probe Time**

Date and time when an instance was last probed.

**Tablespace count**

Number of tablespaces associated with an instance.

**File Count**

Number of data files associated with an instance.

**Log File Count**

Number of log files associated with an instance.

**File Free Space**

Amount of storage free space available on the data files associated with an instance.

**Total Size**

Amount of storage space allocated to an instance.

**File Capacity**

Amount of storage space allocated to the data files associated with an instance.

**By Version****RDBMS Version**

Version of the RDBMS under which an instance is running.

**Computer Name**

Name of the computer on which an instance resides.

**Instance**

Name of an instance.

**RDBMS Type**

RDBMS to which an instance belongs (e.g., Oracle, SQL Server, Sybase, UDB, etc.).

**Port** Port on which an instance listens for requests.

**Home Directory**

Path of the home directory for an instance.

**Logonid**

ID used by Tivoli Storage Productivity Center to log in to an instance.

**Probe Status**

Status of the last probe run against an instance.

**Last Start Time**

Date and time when an instance was last started.

**Last Probe Time**

Date and time when an instance was last probed.

**Tablespace Count**

Number of tablespaces associated with an instance.

**File Count**

Number of data files associated with an instance.

**Log File Count**

Number of log files associated with an instance.

**File Free Space**

Amount of storage free space available on the data files associated with an instance.

**Total Size**

Amount of storage space allocated to an instance.

**File Capacity**

Amount of storage space allocated to the data files associated with an instance.

**By DBMS Type****RDBMS Type**

RDBMS to which an instance belongs (e.g., Oracle, SQL Server, Sybase, UDB, etc.).

**Computer Name**

Name of the computer on which an instance resides.

**Instance**

Name of an instance.

**RDBMS Version**

Version of the RDBMS under which an instance is running.

**Port** Port on which an instance listens for requests.

**Home Directory**

Path of the home directory for an instance.

**Logonid**

ID used by Tivoli Storage Productivity Center to log in to an instance.

**Probe Status**

Status of the last probe run against an instance.

**Last Start Time**

Date and time when an instance was last started.

**Last Probe Time**

Date and time when an instance was last probed.

**Tablespace count**

Number of tablespaces associated with an instance.

**File Count**

Number of data files associated with an instance.

**Log File Count**

Number of log files associated with an instance.

**File Free Space**

Amount of storage free space available on the data files associated with an instance.

**Total Size**

Amount of storage space allocated to an instance.

**File Capacity**

Amount of storage space allocated to the data files associated with an instance.

**By TPC Server****TPC Server**

Name of a subordinate server that is monitoring DBMSs in your environment.

**Total Number of DBMSs**

Number of DBMSs monitored by a subordinate server.

**Number of UDB Instances**

Number of UDB instances monitored by a subordinate server.

**Number of Oracle Instances**

Number of Oracle instances monitored by a subordinate server.

**Number of SQL/Server Instances**

Number of SQL/Server instances monitored by a subordinate server.

**Number of Sybase Instances**

Number of Sybase instances monitored by a subordinate server.

**Tablespace Count**

Number of tablespaces associated with the instances monitored by a subordinate server.

**File Count**

Number of data files associated with the instances monitored by a subordinate server.

**Log File Count**

Number of log files associated with the instances monitored by a subordinate server.

**File Free Space**

Amount of storage free space available on the data files associated with the instances monitored by a subordinate server.

**Total Size**

Amount of storage space allocated to the instances monitored by a subordinate server.

**Total File Capacity (File Capacity)**

Amount of storage space allocated to the data files associated with the instances monitored by a subordinate server.

**Rollup Reports - Database Asset - Oracle, SQL/Server, Sybase, UDB**

Use this report to view information about Oracle, Sybase, SQL Server, and UDB/DB2 instances that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Instance: Shows databases for each computer sorted by computer name.
- By Version: Shows databases for each computer sorted by the version of the RDBMS under which a database is running.
- By TPC Server: Shows databases aggregated at the subordinate server level and sorted by subordinate server name.

**Note:** The first column in each report remains displayed when scrolling through the other columns on the report.

## **By Instance**

### **Computer Name**

Name of the computer on which an Instance resides.

### **Instance**

Name of an instance.

### **RDBMS Type**

RDBMS to which an instance belongs (e.g., Oracle, SQL Server, Sybase, UDB, etc.).

### **RDBMS Version**

Version of the RDBMS under which an instance is running.

**Port** Port on which an instance listens for requests.

### **Home Directory**

Path of the home directory for an instance.

### **Logonid**

ID used by Tivoli Storage Productivity Center to log in to an instance.

### **Probe Status**

Status of the last probe run against an instance.

### **Last Start Time**

Date and time when an instance was last started.

### **Last Probe Time**

Date and time when an instance was last probed.

### **Tablespace count**

Number of tablespaces associated with an instance.

### **File Count**

Number of data files associated with an instance.

### **Log File Count**

Number of log files associated with an instance.

### **File Free Space**

Amount of storage free space available on the data files associated with an instance.

### **Total Size**

Amount of storage space allocated to an instance.

### **File Capacity**

Amount of storage space allocated to the data files associated with an instance.

## **By Version**

### **RDBMS Version**

Version of the RDBMS under which an instance is running.

### **Computer Name**

Name of the computer on which an instance resides.

### **Instance**

Name of an instance.

### **RDBMS Type**

RDBMS to which an instance belongs (e.g., Oracle, SQL Server, Sybase, UDB, etc.).

**Port** Port on which an instance listens for requests.

### **Home Directory**

Path of the home directory for an instance.

### **Logonid**

ID used by Tivoli Storage Productivity Center to log in to an instance.

### **Probe Status**

Status of the last probe run against an instance.

### **Last Start Time**

Date and time when an instance was last started.

### **Last Probe Time**

Date and time when an instance was last probed.

### **Tablespace Count**

Number of tablespaces associated with an instance.

### **File Count**

Number of data files associated with an instance.

### **Log File Count**

Number of log files associated with an instance.

### **File Free Space**

Amount of storage free space available on the data files associated with an instance.

### **Total Size**

Amount of storage space allocated to an instance.

### **File Capacity**

Amount of storage space allocated to the data files associated with an instance.

## **By TPC Server**

### **TPC Server**

Name of a subordinate server that is monitoring DBMSs in your environment.

### **Number of Oracle Instances**

Number of Oracle instances monitored by a subordinate server.

### **Number of SQL/Server Instances**

Number of SQL/Server instances monitored by a subordinate server.

### **Number of Sybase Instances**

Number of Sybase instances monitored by a subordinate server.

**Number of UDB Instances**

Number of UDB instances monitored by a subordinate server.

**Tablespace count**

Number of tablespaces associated with the instances monitored by a subordinate server.

**File Count**

Number of data files associated with the instances monitored by a subordinate server.

**Log File Count**

Number of log files associated with the instances monitored by a subordinate server.

**File Free Space**

Amount of storage free space available on the data files associated with the instances monitored by a subordinate server.

**Total Size**

Amount of storage space allocated to the instances monitored by a subordinate server.

**Total File Capacity**

Amount of storage space allocated to the data files associated with the instances monitored by a subordinate server.

## Rollup Reports - TPC Server Storage Space

Use these reports to view storage metrics related to the disk space, filesystem space, consumed filesystem space, and available filesystem space gathered by the subordinate servers in your environment.

The information shown in these reports is gathered by the master server when it runs "Collecting storage statistics from Tivoli Storage Productivity Center servers (TPC server probes)" on page 340 against associated subordinate servers. If the same storage entity is managed by multiple subordinate servers, rollup reports reflect the storage information from the subordinate server that most recently probed that entity.

### Rollup Reports - TPC Server Storage Space - Disk Space

Use this report to view disk capacity information about the computers and storage systems that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Storage Subsystem: Shows disk capacity aggregated at the storage system level and sorted by storage system name. Storage systems that have been discovered but not probed are not displayed in this report.
- By Storage Subsystem Group: Shows disks capacity aggregated at the storage system monitoring group and storage system reporting group levels and sorted by group name. Storage systems that have been discovered but not probed are not included in the values shown in this report.
- By Computer: Shows disks capacity aggregated at the computer level and sorted by computer name.
- By Computer Group: Shows disks capacity aggregated at the computer monitoring group and computer reporting group levels and sorted by group name.

- By Cluster: Shows disks capacity aggregated at the cluster level and sorted by cluster name.
- By TPC Server (Computers): Shows computer disk capacity aggregated at the subordinate server level and sorted by subordinate server name.
- By TPC Server (Storage Subsystems): Shows storage system disk capacity aggregated at the subordinate server level and sorted by subordinate server name. Storage systems that have been discovered but not probed are not included in the values shown in this report.

**Note:** The first column in each report remains displayed when scrolling through the other columns on the report.

**Note:** The names of some columns have been changed. For your reference, the original names of those columns appear within parentheses ( ) next to their new names.

## By Storage Subsystem

### Subsystem Name

Name of a storage system that is monitored by a subordinate server.

### Subsystem Type

Type of storage system, such as the IBM System Storage SAN Volume Controller.

### Disk Space (Disk Capacity)

Total raw (unformatted) disk capacity of a storage system.

### Available Disk Space (Available Capacity)

Amount of storage space available on the disks associated with a storage system. For Hitachi Data Systems storage systems, internal disk information is not available. This field will be 0.

### Number of Disks

Number of disks associated with a storage system. For Hitachi Data Systems storage systems, internal disk information is not available. This field will be set to 1.

### Number of Volumes (Number of LUNs)

Number of volumes associated with a storage system.

### Storage Subsystem Volume Space (LUN capacity)

Amount of storage space allocated to the volumes associated with the storage system.

### Consumable Storage Subsystem Volume Space (Usable LUN Capacity)

Storage capacity of volumes in the created disk array (not including BCVs, VCM database, and Replication targets).

### Correlated Volume Space (Visible LUN Capacity)

Capacity of all volumes visible to a Storage Resource agent or Data agent associated with a storage system.

### FlashCopy Target Storage Subsystem Volume Space (FlashCopy Target Capacity)

The total capacity of all FlashCopy target volumes in the system. Note: TotalStorage Enterprise Storage Server systems must have at least the following microcode levels: `ess800mincodelevel = 2.4.3.56`, `essf20mincodelevel = 2.3.3.89`. Tivoli Storage Productivity Center does not report FlashCopy information if the TotalStorage Enterprise Storage Server systems do not meet this requirement.

**Volume Space Not Correlated (Not Visible LUN Capacity)**

Usable volume capacity minus assigned volume capacity.

**Unavailable Disk Space (Overhead)**

Amount of capacity that is dedicated to redundancy.

**Formatted Disk Space (Formatted Capacity)**

(IBM TotalStorage Enterprise Storage Server, FAStT, HP StorageWorks, and Hitachi Data Systems only.) Total amount of formatted storage space in the storage pools associated with the storage system. For example, a RAID volume space has more capacity than a stripe-mirror volume space.

**Formatted Disk Space with No Volumes (Total Formatted Space with No Volumes)**

(TotalStorage Enterprise Storage Server, FAStT, HP StorageWorks, and Hitachi Data Systems only.) Amount of storage space available to allocate to volumes in the storage pools associated with the storage system.

**Unformatted Disk Space (Undefined Disk Group or Array Site Capacity)**

Raw disk space in ranks on which no volume space has been defined. For TotalStorage Enterprise Storage Server systems, this value is the total raw capacity of any unformatted disk groups. For IBM XIV Storage System, this value is the total available physical (hard) space that is left for pool creation. For FAStT, HP StorageWorks, and Hitachi Data Systems storage systems, this value is the total capacity of the disks that are not in a storage pool.

**Cache** Amount of RAM or other cache for a storage system.

**Note:** This column is blank for SAN Volume Controllers.

**Assigned Volume Space (Subsystem Assigned LUN Capacity)**

Total storage capacity of all volumes that the storage system knows are assigned to a host computer. Note: The column is blank for SAN Volume Controller.

**Unassigned Volume Space (Subsystem Unassigned LUN Capacity)**

Total storage capacity of all volumes that the storage system knows are not assigned to a host computer. Note: The column is blank for SAN Volume Controller.

**zOS Volume Space (zOS LUN Capacity)**

Total storage capacity for all the volumes allocated to zOS. Data for this column is calculated for TotalStorage Enterprise Storage Server only. The value for this column is set to N/A for other storage systems.

**Open System Volume Space (Open System LUN Capacity)**

Total volume storage capacity (excluding the zOS volume space).

**Physical Disk Space**

Total amount of physical disk space discovered on the monitored storage system arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system disks that become missing after a storage system probe.
- Back-end storage volumes on monitored storage system virtualizers (for example, SAN Volume Controller) that are attached from a storage system.

**Backend Volume Space**

Total amount of storage system volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on



a monitored storage system virtualizer (for example, SAN Volume Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes or virtual storage volumes that are back-end disks on a storage system virtualizer that is not monitored.

## **By Storage Subsystem Group**

### **Subsystem Group Name**

Name of a storage system monitoring group or storage system reporting group.

### **TPC Server**

Name of a subordinate server to which the group is associated.

### **Disk Space (Disk Capacity)**

Total raw (unformatted) disk capacity of the storage systems in a group.

### **Available Disk Space (Available Capacity)**

Amount of storage space available on the disks associated with the storage systems in a group.

### **Number of Disks**

Number of disks associated with the storage systems in a group.

### **Number of Volumes (Number of LUNs)**

Number of volumes associated with the storage systems in a group.

### **Storage Subsystem Volume Space (LUN capacity)**

Amount of storage space allocated to the volumes associated with the storage systems in a group.

### **Consumable Storage Subsystem Volume Space (Usable LUN Capacity)**

Storage capacity of volumes in the created disk array (not including BCVs, VCM database, and Replication targets).

### **FlashCopy Target Storage Subsystem Volume Space (FlashCopy Target Capacity)**

The total capacity of all FlashCopy target volumes in the storage systems in a group.

### **Correlated Volume Space (LUN Capacity Visible to Data Manager)**

Capacity of all volumes visible to a Storage Resource agent or Data agent associated with the storage systems in a group.

### **Volume Space Not Correlated (LUN Capacity Not Visible to Data Manager)**

Usable volume capacity minus assigned volume capacity on the storage systems in a group.

### **Unavailable Volume Space (Overhead)**

Amount of capacity that is dedicated to redundancy in the storage systems in a group.

### **Formatted Space (Formatted Capacity)**

(TotalStorage Enterprise Storage Server, FASTT, HP StorageWorks, and Hitachi Data Systems only.) Total amount of formatted storage space in the storage pools associated with the storage systems in a group.

### **Total Formatted Space with no Volumes (Total Formatted Space with No Volumes)**

(TotalStorage Enterprise Storage Server, FASTT, HP StorageWorks, and

Hitachi Data Systems only.) Amount of storage space available to allocate to volumes in the storage pools associated with the storage systems in a group.

**Unformatted Disk Space (Total Overhead)**

(TotalStorage Enterprise Storage Server only.) Sum of all the space in a disk group minus the formatted capacity on the storage systems in a group.

**Cache** Amount of RAM or other cache for the storage systems in a group.

**Note:** This column does not include values for SAN Volume Controllers.

**Assigned Volume Space (Subsystem Assigned LUN Capacity)**

Total storage capacity of all volumes that the storage systems in a group know are assigned to a host computer.

**Unassigned Volume Space (Subsystem Unassigned LUN Capacity)**

Total storage capacity of all volumes that the storage systems in a group know are not assigned to a host computer.

**zOS Volume Space (zOS LUN Capacity)**

Total storage capacity for all the volumes on the storage systems in a group allocated to zOS.

**Open System Volume Space (Open System LUN Capacity)**

Total volume storage capacity (excluding the zOS volume capacity) of the storage systems in a group.

**Physical Disk Space**

Total amount of physical disk space discovered on the monitored storage system arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system disks that become missing after a storage system probe.
- Back-end storage volumes on monitored storage system virtualizers (for example, SAN Volume Controller) that are attached from a storage system.

**Backend Volume Space**

Total amount of storage system volume space or virtual storage space that Tivoli Storage Productivity Center identifies as the back-end disk space on a monitored storage system virtualizer (for example, SAN Volume Controller). Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system volumes or virtual storage volumes that become missing after a storage system probe.
- Storage system volumes or virtual storage volumes that are back-end disks on a storage system virtualizer that is not monitored.

**By Computer**

**Note:** Two entries for the same host might appear in this report and the values for that host are counted twice in the **TOTAL** column. This scenario occurs under the following circumstances:

- You install two Storage Resource agents on the same host (using the long and short names of the host) that point to different Tivoli Storage Productivity Center servers. For example, you install an agent from serverA using the short name of a host (for example, hostC). You install another agent from serverB using the long name of the same host (for example, hostC.storage.chicago.com). You then set up the servers in a master and subordinate relationship. For

example, serverA is the master and serverB is the subordinate. In this example, there are two entries in the report for the host: hostC and hostC.storage.chicago.com. The values for hostC are counted twice in the **TOTAL** column.

- A Storage Resource agent on a host (hostC) points to two different Tivoli Storage Productivity Center servers (serverA and serverB). The host has a dynamic IP address. You set up serverA as the master and serverB as the subordinate. Server serverA probes hostC, but then the IP address of hostC changes before serverB probes the host. In this example, the report might display two entries for hostC: one entry with the IP address that was probed by serverA; the other entry with the changed IP address that was probed by serverB. The values for hostC are counted twice in the **TOTAL** column.
- A Storage Resource agent on a host (hostC) points to two different Tivoli Storage Productivity Center servers (serverA and serverB). The host has a name that resolves to hostC.storage.chicago.com on serverA and hostC.srm.storage.chicago.com on serverB. You then set up the servers in a master and subordinate relationship. For example, serverA is the master and serverB is the subordinate. In this example, there are two entries in the report for the host: hostC.storage.chicago.com and hostC.srm.storage.chicago.com. The values for hostC are counted twice in the **TOTAL** column.

#### **Computer Name**

Name of a computer.

#### **Disk Space (Total Capacity)**

Total storage capacity for a computer. If the computer is an IBM Tivoli Storage SAN File Systems client, then this value also includes the capacity of any volumes visible to this computer that are owned by the SAN File System.

#### **Available Disk Space (Total Unallocated Space)**

Amount of unused storage space on a computer (not in file systems seen by this operating system). If the computer is an IBM Tivoli Storage SAN File Systems client, then this value excludes the unallocated space on any volumes visible to it that are owned by the SAN File System.

#### **Number of Disks**

Number of disks on a computer.

#### **Owned Disk Space (Owned Capacity)**

Total storage capacity for a computer that is owned by that computer. If Data Manager can determine that a disk was configured for a different host, the capacity for that disk is counted in the "Owned Capacity".

If the computer is an IBM Tivoli Storage SAN File Systems client, then this value does not include the capacity of any volumes visible to it that are owned by the SAN File System. If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and reads the host-ID. If this host-ID does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.

#### **Other Available Disk Space (Other Unallocated Space)**

Total unallocated space on disks owned by the computer. This column does not display the total unallocated disk space on disks not owned by the computer. If the computer is an IBM Tivoli Storage SAN File Systems client, then this value excludes the unallocated space on any volumes visible to it that are owned by the SAN File System.

**Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)**

Total disk capacity of all non-Fibre attached disks for a computer.

**Fibre Channel Attached Disk Space (Total Other Attached Unallocated Space)**

Total unallocated space of all non-Fibre attached disks for a computer.

**Non-Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Capacity)**

Total disk capacity of all fibre channel attached disks for a computer.

**Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Unallocated Space)**

Total unallocated space of all fibre channel attached disks for a computer.

**By Computer Group**

**Computer Group Name**

Name of a computer monitoring group or computer reporting group.

**Disk Space (Total Capacity)**

Total storage capacity for the computers in a group. If the computer is an IBM Tivoli Storage SAN File Systems client, then this value also includes the capacity of any volumes visible to this computer that are owned by the SAN File System.

**Available Disk Space (Total Unallocated Space)**

Amount of unused storage space on the computers in a group.

**Number of Disks**

Number of disks on the computers in a group.

**Owned Disk Space (Owned Capacity)**

Total storage capacity for the computers in a group that is owned by those computers. If Data Manager can determine that a disk was configured for a different host, the capacity for that disk is counted in the "Owned Capacity".

If any computers in the group are IBM Tivoli Storage SAN File Systems clients, then this value does not include the capacity of any volumes visible to it that are owned by the SAN File System. If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and reads the host-ID. If this host-ID does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.

**Other Available Disk Space (Other Unallocated Space)**

Total unallocated space on disks owned by the computers in a group. This column does not display the total unallocated disk space on disks not owned by computers. If any of the computers in the group are IBM Tivoli Storage SAN File Systems clients, then this value excludes the unallocated space on any volumes visible to it that are owned by the SAN File System.

**Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)**

Total disk capacity of all non-Fibre attached disks on the computers in a group.

**Fibre Channel Attached Disk Space (Total Other Attached Unallocated Space)**

Total unallocated space of all non-Fibre attached disks on the computers in a group.

**Non-Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Capacity)**

Total disk capacity of all fibre channel attached disks on the computers in a group.

**Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Unallocated Space)**

Total unallocated space of all fibre channel attached disks on the computers in a group.

**By Cluster**

**Cluster Name**

Name of a cluster.

**Disk Space (Total Capacity)**

Total storage capacity for a cluster.

**Available Disk Space (Total Unallocated Space)**

Amount of unused storage space on a cluster (not in file systems seen by this operating system).

**Owned Disk Space (Owned Capacity)**

Total storage capacity for a cluster that is owned by that cluster. If Data Manager can determine that a disk was configured for a different host, the capacity for that disk is counted in the "Owned Capacity".

If the computer is an IBM Tivoli Storage SAN File Systems client, then this value does not include the capacity of any volumes visible to it that are owned by the SAN File System. If the host is Solaris and Veritas Volume Manager (VxVM) is in use, the product examines the Veritas disk-label and reads the host-ID. If this host-ID does not agree with the local host-ID, the disk is classified as belonging to a different (Solaris) host.

**Other Available Disk Space (Other Unallocated Space)**

Total unallocated space on disks owned by the cluster. This column does not display the total unallocated disk space on disks not owned by the cluster. If a computer in a cluster is an IBM Tivoli Storage SAN File Systems client, then this value excludes the unallocated space on any volumes visible to it that are owned by the SAN File System.

**Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)**

Total disk capacity of all non-Fibre attached disks in a cluster.

**Fibre Channel Attached Disk Space (Total Other Attached Unallocated Space)**

Total unallocated space of all non-Fibre attached disks in a cluster.

**Non-Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Capacity)**

Total disk capacity of all fibre channel attached disks in a cluster.

**Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Unallocated Space)**

Total unallocated space of all fibre channel attached disks in a cluster.

**By TPC Server (Computers)**

**TPC Server**

Name of a subordinate server that is monitoring computer disks in your environment.

**Disk Space (Total Capacity)**

Total storage capacity of all the computer disks monitored by a subordinate server.

**Available Disk Space (Total Unallocated Space)**

Amount of unused storage space on all the computer disks monitored by a subordinate server.

**Number of Disks**

Number of computer disks monitored by a subordinate server.

**Non-Fibre Channel Attached Disk Space (Total Other Attached Capacity)**

Total disk capacity of all non-Fibre attached disks monitored by a subordinate server.

**Fibre Channel Attached Disk Space (Total Other Attached Unallocated Space)**

Total unallocated space of all non-Fibre attached disks monitored by a subordinate server.

**Non-Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Capacity)**

Total disk capacity of all fibre channel attached disks monitored by a subordinate server.

**Fibre Channel Attached Available Disk Space (Total Fibre Channel Attached Unallocated Space)**

Total unallocated space of all fibre channel attached disks monitored by a subordinate server.

**By TPC Server (Storage Subsystems)****TPC Server**

Name of a subordinate server that is monitoring storage system disks in your environment.

**Disk Space (Total Capacity)**

Total raw (unformatted) disk capacity of the storage systems monitored by a subordinate server.

**Available Disk Space (Total Unallocated Space)**

Amount of storage space available on the disks associated with the storage systems monitored by a subordinate server.

**Number of Disks**

Number of storage system disks monitored by a subordinate server.

**Physical Disk Space**

Total amount of physical disk space discovered on the monitored storage system arrays. Tivoli Storage Productivity Center does not include the following volumes in its calculation of the value for this column:

- Storage system disks that become missing after a storage system probe.
- Back-end storage volumes on monitored storage system virtualizers (for example, SAN Volume Controller) that are attached from a storage system.

**Rollup Reports - TPC Server Storage Space - File System Space**

Use this report to view file system space information about the file systems that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Computer: Shows file system space aggregated at the computer level and sorted by computer name.
- By Computer Group: Shows file system space aggregated at the computer monitoring group and computer reporting group levels and sorted by group name.
- By Cluster: Shows file system space aggregated at the cluster level and sorted by cluster name.
- By TPC Server: Shows file system space aggregated at the subordinate server level and sorted by subordinate server name.

**Note:** The first column in each report remains displayed when scrolling through the other columns on the report.

## By Computer

### Computer Name

Name of the computer on which a file system resides.

### File System Space

Total amount (capacity) of storage space in the file systems on a computer. The value in this column might be greater than **Consumed File System Space + Available File System Space** for some UNIX systems. This situation occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the **df** command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **File System Space** column is greater than or equal to the value of **Consumed File System Space + Available File System Space**, the difference representing the space wasted due to the system overhead.

### Consumed File System Space

Total amount of used storage space in the file systems on a computer. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead). Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

The value that appears in the **TOTAL** row for this column represents the difference between the **TOTAL** values in the **File System Space** and **Available File System Space** columns. If the calculated value for the **TOTAL** row is negative, the column displays 0.

The **TOTAL** row does not add the used space coming from VMWare computers. See the *VMWare capacity reports* topic in the *Planning* chapter of the *Tivoli Storage Productivity Center Installation and Configuration Guide* for information about how file system and disk capacity are calculated in this report for VMWare computers. This topic is also available in the *Planning* section of the *Tivoli Storage Productivity Center Information Center*.

The value that appears in the **TOTAL** row for this column does not include the following file systems:

- File systems on discovered hypervisors that are assigned to virtual machine servers where the Storage Resource agent or Data agent has been installed.
- File systems on discovered virtual machine servers where the Storage Resource agent or Data agent is installed, but the hypervisor for the virtual machine servers has not been discovered and the consumed file system space exceeds the capacity of the total file system space.

#### **Available File System Space**

Total amount of unused storage space in the file systems on a computer.

#### **Number of FS's/LV's**

Number of file systems or logical volumes on a computer.

### **By Computer Group**

#### **Computer Group**

Name of a computer monitoring group or computer reporting group.

#### **File System Space**

Total amount (capacity) of storage space on the file systems in a group. Note that the value in this column might be greater than **Consumed File System Space + Available File System Space** for some UNIX systems. This situation occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the **df** command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **File System Space** column is greater than or equal to the value of **Consumed File System Space + Available File System Space**, the difference representing the space wasted due to the system overhead.

#### **Consumed File System Space**

Total amount of used storage space on the file systems in a group. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

#### **Available File System Space**

Total amount of unused storage space on the file systems in a group.

#### **Number of FS's/LV's**

Number of file systems or logical volumes in a group.

### **By Cluster**

#### **Cluster**

Name of a cluster.

#### **File System Space**

Total storage capacity of the file systems in a cluster. The value in this column might be greater than **Consumed File System Space + Available File System Space** for some UNIX systems. This situation occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the **df** command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **File System Space** column is greater than or equal to the value of **Consumed File System Space + Available File System Space**, the difference representing the space wasted due to the system overhead.



**Consumed File System Space**

Total amount of used storage space on the file systems in a cluster. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

**Available File System Space**

Total amount of unused storage space on the file systems in a cluster.

**By TPC Server****TPC Server**

Name of a subordinate server that is monitoring file systems in your environment.

**File System Space**

Total storage capacity of all the file systems monitored by a subordinate server.

**Consumed File System Space**

Total amount of used storage space on the file systems monitored by a subordinate server. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

**Available File System Space**

Total amount of unused storage space on the file systems monitored by a subordinate server.

**Total Number of FS's/LV's**

Total number of file systems or logical volumes monitored by a subordinate server.

**Rollup Reports - TPC Server Storage Space - Consumed File System Space**

Use this report to view information about the consumed storage space on file systems that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Computer: Shows file system consumed space aggregated at the computer level and sorted by computer name.
- By Computer Group: Shows file system consumed space aggregated at the computer monitoring group and computer reporting group levels and sorted by group name.
- By Cluster: Shows file system consumed space aggregated at the cluster level and sorted by cluster name.
- By TPC Server: Shows file system consumed space aggregated at the subordinate server level and sorted by subordinate server name.

**Note:** The first column in each report remains displayed when scrolling through the other columns on the report.

**By Computer****Computer Name**

Name of the computer on which a file system resides.

**Consumed File System Space**

Total amount of used storage space in the file systems on a computer. On some UNIX file systems, this value does not include storage space reserved

by the operating system (overhead). Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

The value that appears in the **TOTAL** row for this column represents the difference between the **TOTAL** values in the **File System Space** and **Available File System Space** columns. If the calculated value for the **TOTAL** row is negative, the column displays 0.

The **TOTAL** row does not add the used space coming from VMWare computers. See the *VMWare capacity reports* topic in the *Planning* chapter of the *Tivoli Storage Productivity Center Installation and Configuration Guide* for information about how file system and disk capacity are calculated in this report for VMWare computers. This topic is also available in the *Planning* section of the *Tivoli Storage Productivity Center Information Center*.

The value that appears in the **TOTAL** row for this column does not include the following file systems:

- File systems on discovered hypervisors that are assigned to virtual machine servers where the Storage Resource agent or Data agent has been installed.
- File systems on discovered virtual machine servers where the Storage Resource agent or Data agent is installed, but the hypervisor for the virtual machine servers has not been discovered and the consumed file system space exceeds the capacity of the total file system space.

#### **Available File System Space**

Total amount of unused storage space in the file systems on a computer.

#### **File System Space**

Total amount (capacity) of storage space in the file systems on a computer. The value in this column might be greater than **Consumed File System Space** + **Available File System Space** for some UNIX systems. This situation occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the **df** command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **File System Space** column is greater than or equal to the value of **Consumed File System Space** + **Available File System Space**, the difference representing the space wasted due to the system overhead.

#### **Number of FS's/LV's**

Number of file systems or logical volumes on a computer.

#### **By Computer Group**

##### **Computer Group**

Name of a computer monitoring group or computer reporting group.

##### **File System Space**

Total amount (capacity) of storage space on the file systems in a group. The value in this column might be greater than **Consumed File System Space** + **Available File System Space** for some UNIX systems. This situation occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the **df**

command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **File System Space** column is greater than or equal to the value of **Consumed File System Space** + **Available File System Space**, the difference representing the space wasted due to the system overhead.

#### **Consumed File System Space**

Total amount of used storage space on the file systems in a group. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

#### **Available File System Space**

Total amount of unused storage space on the file systems in a group.

#### **Number of FS's/LV's**

Number of file systems or logical volumes in a computer group.

### **By Cluster**

#### **Cluster**

Name of a cluster.

#### **Consumed File System Space**

Total amount of used storage space on the file systems in a cluster. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

#### **Available File System Space**

Total amount of unused storage space on the file systems in a cluster.

#### **File System Space**

Total storage capacity of the file systems in a cluster. The value in this column might be greater than **Consumed File System Space** + **Available File System Space** for some UNIX systems. This situation occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the **df** command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **File System Space** column is greater than or equal to the value of **Consumed File System Space** + **Available File System Space**, the difference representing the space wasted due to the system overhead.

#### **Number of FS's/LV's**

Number of file systems or logical volumes in a cluster.

### **By TPC Server**

#### **TPC Server**

Name of a subordinate server that is monitoring file systems in your environment.

#### **Consumed File System Space**

Total amount of used storage space on the file systems monitored by a subordinate server. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

#### **Available File System Space**

Total amount of unused storage space on the file systems monitored by a subordinate server.

#### **File System Space**

Total storage capacity of all the file systems monitored by a subordinate server.

### Total Number of FS's/LV's

Number of file systems or logical volumes monitored by a subordinate server.

## Rollup Reports - TPC Server Storage Space - Available File System Space

Use this report to view information about the available storage space on file systems that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Computer: Shows available file system space aggregated at the computer level and sorted by computer name.
- By Computer Group: Shows available file system space aggregated at the computer monitoring group and computer reporting group levels and sorted by group name.
- By Cluster: Shows available file system space aggregated at the cluster level and sorted by cluster name.
- By IBM Tivoli Storage Productivity Center Server: Shows available file system space aggregated at the subordinate server level and sorted by subordinate server name.

**Note:** The first column in each report remains displayed when scrolling through the other columns on the report.

### By Computer

#### Computer Name

Name of the computer on which a file system resides.

#### Available File System Space

Total amount of unused storage space in the file systems on a computer.

#### Consumed File System Space

Total amount of used storage space in the file systems on a computer. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead). Tivoli Storage Productivity Center does not include the following file systems in its calculation of the value for this column:

- File systems that become missing after a probe.
- File systems that are mounted remotely and are not identified to a network attached storage server.
- File systems that are not mounted.

The value that appears in the **TOTAL** row for this column represents the difference between the **TOTAL** values in the **File System Space** and **Available File System Space** columns. If the calculated value for the **TOTAL** row is negative, the column displays 0.

The **TOTAL** row does not add the used space coming from VMWare computers. See the *VMWare capacity reports* topic in the *Planning* chapter of the *Tivoli Storage Productivity Center Installation and Configuration Guide* for information about how file system and disk capacity are calculated in this report for VMWare computers. This topic is also available in the *Planning* section of the *Tivoli Storage Productivity Center Information Center*.

The value that appears in the **TOTAL** row for this column does not include the following file systems:

- File systems on discovered hypervisors that are assigned to virtual machine servers where the Storage Resource agent or Data agent has been installed.
- File systems on discovered virtual machine servers where the Storage Resource agent or Data agent is installed, but the hypervisor for the virtual machine servers has not been discovered and the consumed file system space exceeds the capacity of the total file system space.

### **File System Space**

Total amount (capacity) of storage space in the file systems on a computer. The value in this column might be greater than **Consumed File System Space + Available File System Space** for some UNIX systems. This situation occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the **df** command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **File System Space** column is greater than or equal to the value of **Consumed File System Space + Available File System Space**, the difference representing the space wasted due to the system overhead.

### **Number of FS's/LV's**

Number of file systems or logical volumes on a computer.

## **By Computer Group**

### **Computer Group**

Name of a computer monitoring group or computer reporting group.

### **Available File System Space**

Total amount of unused storage space on the file systems in a group.

### **Consumed File System Space**

Total amount of used storage space on the file systems in a group. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

### **File System Space**

Total amount (capacity) of storage space on the file systems in a group. The value in this column might be greater than **Consumed File System Space + Available File System Space** for some UNIX systems. This situation occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the **df** command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **File System Space** column is greater than or equal to the value of **Consumed File System Space + Available File System Space**, the difference representing the space wasted due to the system overhead.

### **Number of FS's/LV's**

Number of file systems or logical volumes in a computer group.

## **By Cluster**

### **Cluster**

Name of a cluster.

### **Available File System Space**

Total amount of unused storage space on the file systems in a cluster.

**Consumed File System Space**

Total amount of used storage space on the file systems in a cluster. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

**File System Space**

Total storage capacity of the file systems in a cluster. The value in this column might be greater than **Consumed File System Space + Available File System Space** for some UNIX systems. This situation occurs because the used space reported on UNIX (both by the APIs used by the Storage Resource agents or Data agents and by the **df** command) does not include some space reserved by the operating system (overhead). Therefore, the value in the **File System Space** column is greater than or equal to the value of **Consumed File System Space + Available File System Space**, the difference representing the space wasted due to the system overhead.

**By TPC Server****TPC Server**

Name of a subordinate server that is monitoring file systems in your environment.

**Available File System Space**

Total amount of unused storage space on the file systems monitored by a subordinate server.

**Consumed File System Space**

Total amount of used storage space on the file systems monitored by a subordinate server. On some UNIX file systems, this value does not include storage space reserved by the operating system (overhead).

**File System Space**

Total storage capacity of all the file systems monitored by a subordinate server.

**Number of FS's/LV's**

Number of file systems or logical volumes monitored by a subordinate server.

## Rollup Reports - Database Capacity

Use these reports to view storage capacity, used space, and free space at the computer, computer group, database, instance, and IBM Tivoli Storage Productivity Center server level for RDBMSs that are monitored by subordinate servers.

For a specific instance, computer, group of computers, or database that is monitored by a subordinate server, you can use these rollup reports to view:

- how much storage you have for an RDBMS and the objects it contains
- how much storage is being used for an RDBMS and the objects it contains
- what percentage of RDBMS storage is free for expansion

The information shown in these reports is gathered by the master server when it runs Tivoli Storage Productivity Center server probes against associated subordinate servers. If the same storage entity is managed by multiple subordinate servers, rollup reports reflect the storage information from the subordinate server that most recently probed that entity.

## Rollup Reports - Database Capacity - All DBMSs

Use this report to view storage capacity information for the DBMSs (including Oracle, Sybase, SQL Server, and UDB/DB2) that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- **By Instance:** Shows database capacity aggregated at the computer level and sorted by instance name.
- **By Computer:** Shows database capacity aggregated at the computer level and sorted by computer name.
- **By Computer Group:** Shows database capacity aggregated at the computer monitoring group and computer reporting group levels and sorted by group name.
- **By TPC Server:** Shows database capacity aggregated at the subordinate server level and sorted by subordinate server name.

**Note:** The first column in each report remains displayed when scrolling through the other columns on the report.

### By Instance

#### Instance

Name of an instance.

#### RDBMS Type

RDBMS to which an instance belongs (e.g., Oracle, SQL Server, Sybase, UDB, etc.).

#### Computer Name

Name of the computer on which an Instance resides.

#### Total Size

Amount of storage space allocated to an instance.

#### File Capacity

Amount of storage space allocated to the data files associated with an instance.

#### File Free Space

Amount of storage free space available on the data files associated with an instance.

#### Tablespace Count

Number of tablespaces associated with an instance.

#### File Count

Number of data files associated with an instance.

#### Log File Count

Number of log files associated with an instance.

### By Computer

#### Computer Name

Name of the computer on which an Instance resides.

#### Total Size

Amount of storage space allocated to an instance.

**File Capacity**

Amount of storage space allocated to the data files associated with an instance.

**File Free Space**

Amount of storage free space available on the data files associated with an instance.

**Tablespace Count**

Number of Tablespaces associated with an instance.

**File Count**

Number of data files associated with an instance.

**Log File Count**

Number of log files associated with an instance.

**By Computer Group****Computer Group**

Name of a computer monitoring group or reporting group.

**Total Size**

Amount of storage space allocated to instances in a group.

**File Capacity**

Amount of storage space allocated to the data files associated with instances in a group.

**File Free Space**

Amount of storage free space available on the data files associated with instances in a group.

**Tablespace Count**

Number of databases or tablespaces associated with instances in a group.

**File Count**

Number of data files associated with instances in a group.

**Log File Count**

Number of log files associated with instances in a group.

**TPC Server**

Name of a subordinate server that is monitoring the computer group.

**By TPC Server****TPC Server**

Name of a subordinate server that is monitoring DBMSs in your environment.

**Total Size**

Amount of storage space allocated to instances monitored by a subordinate server.

**Total File Capacity**

Amount of storage space allocated to the data files associated with instances monitored by a subordinate server.

**File Free Space**

Amount of storage free space available on the data files associated with instances monitored by a subordinate server.



**Tablespace Count**

Number of tablespaces associated with instances monitored by a subordinate server.

**File Count**

Number of data files associated with instances monitored by a subordinate server.

**Log File Count**

Number of log files associated with instances monitored by a subordinate server.

**Rollup Reports - Database Capacity - Oracle, Sybase, SQL Server, and UDB**

Use this report to view storage capacity information for the Oracle, Sybase, SQL Server, and UDB/DB2 databases that are monitored by subordinate servers in your environment.

You can generate and sort the data within this report using the following by clauses:

- By Database: Shows database capacity aggregated at the computer level and sorted by instance name.
- By Computer: Shows database capacity aggregated at the computer level and sorted by computer name.
- By Computer Group: Shows database capacity aggregated at the computer monitoring group and computer reporting group levels and sorted by group name.
- By TPC Server: Shows database capacity aggregated at the subordinate server level and sorted by subordinate server name.

**Note:** The first column in each report remains displayed when scrolling through the other columns on the report.

**By Database****Database Instance**

Name of a database instance.

**Computer Name**

Name of the computer on which an instance resides.

**Total Size**

Amount of storage space allocated to an instance.

**File Capacity**

Amount of storage space allocated to the data files associated with an instance.

**File Free Space**

Amount of storage free space available on the data files associated with an instance.

**Tablespace Count**

Number of tablespaces associated with an instance.

**File Count**

Number of data files associated with an instance.

**Log File Count**

Number of log files associated with an instance.

## **By Computer**

### **Computer Name**

Name of the computer on which an instance resides.

### **Total Size**

Amount of storage space allocated to an instance.

### **File Capacity**

Amount of storage space allocated to the data files associated with an instance.

### **File Free Space**

Amount of storage free space available on the data files associated with an instance.

### **Tablespace Count**

Number of tablespaces associated with an instance.

### **File Count**

Number of data files associated with an instance.

### **Log File Count**

Number of log files associated with an instance.

## **By Computer Group**

### **Computer Group**

Name of a computer monitoring group or reporting group.

### **Total Size**

Amount of storage space allocated to instances in a group.

### **File Capacity**

Amount of storage space allocated to the data files associated with instances in a group.

### **File Free Space**

Amount of storage free space available on the data files associated with instances in a group.

### **Tablespace Count**

Number of tablespaces associated with instances in a group.

### **File Count**

Number of data files associated with instances in a group.

### **Log File Count**

Number of log files associated with instances in a group.

### **TPC Server**

Name of a subordinate server that is monitoring the group.

## **By TPC Server**

### **TPC Server**

Name of a subordinate server that is monitoring DBMSs in your environment.

### **Total Size**

Amount of storage space allocated to instances monitored by a subordinate server.

**Total File Capacity**

Amount of storage space allocated to the data files associated with instances monitored by a subordinate server.

**File Free Space**

Amount of storage free space available on the data files associated with instances monitored by a subordinate server.

**Tablespace Count**

Number of tablespaces associated with instances monitored by a subordinate server.

**File Count**

Number of data files associated with instances monitored by a subordinate server.

**Log File Count**

Number of log files associated with instances monitored by a subordinate server.

---

## Data source reports - overview

Use data source reports to view detailed information about the agents from which IBM Tivoli Storage Productivity Center collects information about storage resources. Data sources include CIMOM agents, Storage Resource agents, Data agents, Fabric agents, VMWare, and subordinate Tivoli Storage Productivity Center servers.

### Data source reports - CIMOM agents

Use these reports to view detailed information about the CIMOM agents that have been associated with IBM Tivoli Storage Productivity Center.

**Related Topics**

- Selecting resources upon which to report
- Filtering the resources that appear in a report
- Generating a report
- Drilling down a report row (By Managed Device report only)
- Navigating tabs

**Data source report - By CIMOM Agent**

Use this report to view detailed information about the CIMOM agents that are used as data sources for IBM Tivoli Storage Productivity Center.

Each CIMOM that has been associated with IBM Tivoli Storage Productivity Center appears as its own row in this report.

**Service URL**

The service URL of the CIMOM containing the IP address of the CIMOM, the port on which the CIMOM is listening, and the protocol used for communication. This URL has a protocol [http|https], an IP address (IPv4 or IPv6), and a port number. For example:

- **IPv4:** https://127.0.0.1:5989
- **IPv6:** https://[2001:DB8::1234:0000:0000:5678:ABCD]:5989

**Display Name**

The name of the CIMOM as specified by the CIMOM provider that will appear in the Tivoli Storage Productivity Center user interface.

**Description**

The optional description that was entered on the Add CIMOM panel.

**Managed Device Category**

The type of device that is managed by a CIMOM (for example, storage subsystem, switch, or tape library). The value **NA** appears in this column if you have not run a CIMOM discovery against a CIMOM or a CIMOM does not manage any devices.

**Connection Status**

The status of this CIMOM with respect to Tivoli Storage Productivity Center. Possible values are: **SUCCESS**, **UNCONFIGURED**, **UNKNOWN**, **INVALID\_NAMESPACE**, **TIMEOUT**, **REFUSED**, **LOGIN\_FAILED**, **SSL\_HANDSHAKE\_ERROR**, **SSL\_REGISTRATION\_INVALID**, **CIMCLIENT\_ERROR**

**Status Timestamp**

The date and time when the **Connection Status** information was last collected.

**Number of Managed Devices**

The number of devices managed by a CIMOM.

**Interoperability Namespace**

The interoperability namespace of a CIMOM. This namespace within the CIMOM allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how Tivoli Storage Productivity Center interacts with the CIMOM when retrieving information. For a NetApp CIMOM, enter `/interop`.

**Username**

The CIMOM user name used for authentication.

**Truststore Location**

The location (path on this computer) of a certificate file for certificate based authentication in the https protocol.

**User Interface Description**

The name of the Human Interface Service (if any) supported by this CIMOM.

**Software Level**

The software version level of the CIMOM agent.

**Protocol Version**

The version of the cim-xml protocol.

**Authentication Mechanism**

The authentication mechanism supported by the CIMOM. This column can contain the following values: **Unknown**, **None**, **Other**, **Basic**, **Digest**.

**Alias** The alias of the CIMOM.

**Protocol**

The protocol used to communicate with the CIMOM. Possible values are **http** and **https**.

**SLP Attributes**

The standard set of attributes for this CIMOM. The attributes are retrieved via SLP.

## Data source report - By Managed Device

Use this report to view information about the devices managed by CIMOMs in your environment. Only devices that are managed by the CIMOMs used as data sources for IBM Tivoli Storage Productivity Center are displayed.

You must run a CIMOM discovery job before you can view the devices managed by CIMOMs in this report. If a device is managed by multiple CIMOMs, the report displays a separate row for each of those CIMOMs. Therefore, the number of rows for a device in the report is equal to the number of CIMOMs managing that device.

### Magnifying glass icon

Click this button to view a more detailed report about a device.

#### Device Name

The display name (label) of the managed storage subsystem, switch, or tape library.

#### Device Category

Indicates whether a device that is managed by a CIMOM is a storage subsystem, switch, or tape library.

#### Device Type

Indicates the type of a managed device based on its category. For example, a device that is a storage subsystem might have a value of **DS8000** in this column.

#### Device Status

The consolidated status of the managed device.

#### Device Operational Status

The operational status of a storage subsystem.

#### Device Last Probe Time

The date and time when a probe was last run on a managed device.

#### Device Vendor

The manufacturer of the managed device.

#### Device Model

The model number of the managed device.

#### Device Serial Number

The serial number of the managed device.

#### Device Microcode Level

The microcode or firmware level of the managed device.

#### Device IP Address

The IP address of the managed device.

#### Device Element Manager

The URL or the fully-qualified name of the executable file that launches the element manager.

#### CIMOM Service URL

The service URL of the CIMOM. This URL includes the IP address or host name of the CIMOM, the port on which the CIMOM is listening, and the protocol used for communication (http | https). For example: `https://127.0.0.1:9551`.

**CIMOM Display Name**

The name of the CIMOM as specified by the CIMOM provider that will appear in the Tivoli Storage Productivity Center user interface.

**CIMOM Description**

The optional description that was entered on the Add CIMOM panel.

**CIMOM Connection Status**

The status of this CIMOM with respect to Tivoli Storage Productivity Center. Possible values are: SUCCESS, UNCONFIGURED, UNKNOWN, INVALID\_NAMESPACE, TIMEOUT, REFUSED, LOGIN\_FAILED, SSL\_HANDSHAKE\_ERROR, SSL\_REGISTRATION\_INVALID, CIMCLIENT\_ERROR

**CIMOM Status Timestamp**

The date and time when the **Connection Status** information was last collected.

**CIMOM Interoperability Namespace**

The interoperability namespace of a CIMOM. This namespace within the CIMOM allows for accessing the CIM Interop Schema (including the class instances of the Server Profile) and determines how Tivoli Storage Productivity Center interacts with the CIMOM when retrieving information. For a NetApp CIMOM, enter /interop.

**CIMOM Username**

The CIMOM user name used for authentication.

**CIMOM Trustore Location**

The location (path on this computer) of a certificate file for certificate based authentication in the https protocol.

**CIMOM User Interface Description**

The name of the Human Interface Service (if any) supported by this CIMOM.

**CIMOM Software Level**

The software version level of the CIMOM agent.

**CIMOM Protocol Version**

The version of the cim-xml protocol.

**CIMOM Authentication Mechanism**

The authentication mechanism supported by the CIMOM. This column can contain the following values: **Unknown**, **None**, **Other**, **Basic**, **Digest**.

**CIMOM Alias**

The alias of the CIMOM.

**CIMOM Protocol**

The protocol used to communicate with the CIMOM. Possible values are **http** and **https**.

**CIMOM SLP Attributes**

The standard set of attributes for this CIMOM. The attributes are retrieved via SLP.

---

## Storage tier reports - overview

Use storage tier reports to help you determine the best storage tier configuration for the needs of your environment. These reports can identify storage that is not allocated appropriately for the applications that use that storage. You can use the data that IBM Tivoli Storage Productivity Center collects from IBM System Storage SAN Volume Controller, IBM Storwize V7000 Unified, and IBM Storwize V7000 managed-disk groups (MDisk groups) to create and run storage tier performance reports on those managed-disk groups and related VDisks. You can set performance and usage thresholds for the managed-disk groups and VDisk, and the storage tier reports indicate when those thresholds are exceeded for those devices and entities. Tivoli Storage Productivity Center provides Tivoli Common Reporting reports that show performance and usage data for the IBM System Storage SAN Volume Controller, IBM Storwize V7000 Unified, and IBM Storwize V7000 systems in your enterprise.

### Running storage tier reports

Follow this procedure to run storage tier reports on managed disk (MDisk) groups, volumes (VDisks) volumes, and application storage usage and activity using Tivoli Common Reporting.

#### Storage tier reports

Storage tier reports help you determine whether your storage is tiered correctly for the needs of your environment. These reports can identify storage that is not allocated appropriately for the applications that are using that storage. Tiered storage is the assignment of different categories of data to different types of storage media to reduce total storage cost. For example, a data repository might require high I/O capability when it is first implemented. As the repository grows larger, there might be portions of that data that are not accessed frequently. Based on changing access and usage trends, some of the data on the storage pools used by this repository can be migrated to lower-tier storage.

Other storage devices can require a faster read/write response or increased space requirements based on the applications that are using those devices. Moving those repositories to storage with a faster read/write response time and larger capacity can help balance the storage and network activity.

Storage tier reports provide information that can help you make storage-based performance decisions:

- Balancing workloads between MDisk groups.
- Identifying MDisk groups that satisfy specific performance and capacity requirements.
- Setting thresholds related to storage tiering, and modifying your storage tiering based on these thresholds.
- Adding new applications based on their storage and performance requirements.

Use the data that Tivoli Storage Productivity Center collects from SAN Volume Controller and Storwize V7000 MDisk groups to run storage tier performance reports on those MDisk groups and related VDisks. You can set performance and usage thresholds for the MDisk groups and VDisks. The storage tier reports indicate when those thresholds are exceeded for those devices and entities. Tivoli

Storage Productivity Center provides Tivoli Common Reporting reports that show performance and usage data for the SAN Volume Controller and Storwize V7000 systems in your enterprise.

Use the information displayed in the storage tier reports provided by Tivoli Storage Productivity Center to help you improve your storage utilization. The storage tier reports are:

- MDisk Groups - Workload Activity Report
- MDisk Group Details Report
- MDisk Groups - VDisk Workload Activity Report
- MDisk Group - VDisk Workload Activity Details Report
- Storage Virtualizer - VDisk Workload Activity Report
- Storage Resource Group - VDisk Workload Activity Report
- Storage Resource Group - VDisk Workload Activity Details Report
- VDisk Details Report

## Configuring Tivoli Storage Productivity Center for storage tier reports

Use the following steps to prepare Tivoli Storage Productivity Center to run storage tier reports. You can then install and configure Tivoli Common Reporting and run the storage tier reports.

**Note:** You can install and use the Tivoli Common Reporting functions for other purposes in addition to the storage tier reports.

1. Add a storage device and gather information about that device.
  - Use the Configure Devices wizard to set up storage devices for monitoring by Tivoli Storage Productivity Center. The wizard guides you through the following steps:
    - Adding a device as a data source
    - Running a discovery, including devices in groups
    - Specifying alerts
    - Setting up data collection schedules
  - After saving the storage device configuration, Tivoli Storage Productivity Center automatically runs a probe job to gather information about the new device.
2. Extract and automatically set the back-end pool characteristics for the device based on the naming convention used for the storage pools. The naming convention that you use for pool names must contain the following information:
  - Back-end type of storage system, RAID, or disk type
  - Number of disks used by the storage pool

Use a default naming pattern or create a custom naming pattern for the storage pools. Use the **autosetarray** command to extract information from the pool name in storage device and to set these values and the number of disks that the pool uses.

The following command extracts information about pools that match the default pattern in storage subsystem 00000200A0C0005C+0:

```
tpctool> autosetarray -subsysid 00000200A0C0005C+0
```

The following error messages and output are displayed:



[AAJ002013E] The pool name mdiskgroup0 is too short to be matched against the pattern BxxxxxxxDRDDx\*.  
 [AAJ002013E] The pool name Cognos is too short to be matched against the pattern BxxxxxxxDRDDx\*.  
 [AAJ002013E] The pool name mdiskgrp1 is too short to be matched against the pattern BxxxxxxxDRDDx\*.  
 [AAJ002013E] The pool name mdiskgrp2 is too short to be matched against the pattern BxxxxxxxDRDDx\*.

| Name of Array | Back-end | Type | Back-end RAID Type | Back-end Disk Type |
|---------------|----------|------|--------------------|--------------------|
| mdiskgroup0   | -        | -    | -                  | -                  |
| Cognos        | -        | -    | -                  | -                  |
| mdiskgrp1     | -        | -    | -                  | -                  |
| mdiskgrp2     | -        | -    | -                  | -                  |
| Cpool3GA60713 | C        | 6    |                    | A07                |
| Cpool2GA607   | C        | 6    |                    | A07                |
| Dpool1GA607F  | D        | 6    |                    | A07                |

| Back-end | Disk Count | Status     |
|----------|------------|------------|
| -        |            | No Match   |
| -        |            | No Match   |
| -        |            | No Match   |
| -        |            | No Match   |
| 1        |            | Successful |
| 1        |            | Successful |
| 1        |            | Successful |

For more information about the **autosetarray** command, see the IBM Tivoli Storage Productivity Center Command-line Interface Reference. For information about the **tpctool** command, see the instructions about using the CLI to set the back-end pool characteristics of an MDisk group.

- Set the back-end pool characteristics for each MDisk group on which you want to run Tivoli Common Reporting reports:

**Restriction:** You cannot set pool characteristics for MDisk group using active easy tiering, or that use mixed disks.

In the GUI, do the following:

- Click **Disk Manager > Storage Subsystems > select a storage subsystem > MDisk Groups**.
- Click the icon next to the MDisk group name in the list to display the **Details** page.
- Set the back-end pool characteristics for the MDisk group by using the fields in the Information about the Back-end Subsystem Pool section.
- 

**Type** EMC Clarion, EMC Symmetrix, DS8000, DS4000, and XIV

**RAID type**

RAID 1, RAID 5, RAID 6, RAID 10, and RAID X

**Disk type**

Fiber - 15,000 rpm, Fiber - 10,000 rpm, and SATA 7,500 rpm

**Number of physical disks**

A number greater than zero

- Click **Save**.

Read I/O capability is calculated when all the preceding attributes are entered. This measurement is I/O per second.

In the CLI, select one of these options: :

- a. Start an interactive CLI session by using one of the interactive CLI sessions
  - with or without credentials.
  - Interactive CLI session:
    - 1) Open a command prompt window.
    - 2) Go to `install directory\cli\tpctool` directory.
    - 3) From the `tpctool>` command prompt, type **tpctool** and press Enter.
    - 4) Enter CLI commands at this prompt.
    - 5) Enter a value for the Tivoli Storage Productivity Center user, pwd, and url parameters for each command.
  - Interactive CLI session with credentials:
    - 1) Type the following command at a system prompt using the Tivoli Storage Productivity Center user name and password:  
`tpctool -user <user_name> -pwd <user_password> -url <Device_server_URL>`
    - 2) From the `tpctool>` command prompt, enter CLI commands. You are not required to enter the `tpctool` keyword and user, pwd, and url for each command.
- b. To identify the subsystem device and display the device attributes Use the **lsarray** command.
  - For example: `tpctool> lsarray -dev 0000020064405BA0+0 -1`
  - This displays device attributes, such as:
    - BE Read I/O Cap
    - BE Type, BE RAID Type
    - BE Disk Type
    - BE Disk Count

**Note:** The read I/O capability value defaults to 0. If the specified pool does not have values for the other back-end attributes, the command output displays a -.

- c. To specify values for the subsystem device, use the **setarray** command.
  - For example:
 

```
tpctool> setarray -array 0000020064405BA0:0+0000020064405BA0+0
-backendtype D -backendraidthe 5
-backenddisktype F10 -backenddiskcount 160
```
  - This sets the following device attributes: BE Type, BE RAID Type, BE Disk Type, and BE Disk Count.

Command output:

```
Array List Status
=====
0000020064405BA0:0+0000020064405BA0+0 Succeeded.
```

You must specify at least one of these attributes. Values for unspecified attributes are left unchanged if that storage pool attribute already has a value. If an attribute has no value, the appropriate DEFAULT value is used.

**Note:** The read I/O capability value defaults to 0. If the specified pool does not have values for the other back-end attributes, the command output shows a -.

- d. (Optional) Update the alert threshold values for the storage pools by using the **getdscf** and **setdscfg** commands.
  - Use the **getdscfg** command to display the MDisk group attribute:  
 The example is `tpctool> getdscfg -context ManagedDiskGroup`  
 The output is:

| Property                           | Context          | Value |
|------------------------------------|------------------|-------|
| ManagedDiskGroupAvailSize          | ManagedDiskGroup | 80    |
| ManagedDiskGroupAvailIORate        | ManagedDiskGroup | 80    |
| ManagedDiskGroupBEIORExceeds       | ManagedDiskGroup | 80    |
| ManagedDiskGroupBEIOWExceeds       | ManagedDiskGroup | 80    |
| ManagedDiskGroupBERResponseExceeds | ManagedDiskGroup | 80    |
| ManagedDiskGroupBEWResponseExceeds | ManagedDiskGroup | 80    |

Alerts are generated for the following conditions:

- The threshold for Pool Available Size exceeds 80% of total available size.
- The threshold for Pool Available IO Rate exceeds 80% of the available capability.
- The maximum amount of time during Pool BE I/O Read Rate exceeds 80% of the actual maximum.
- The maximum amount of time during Pool BE I/O Write Rate exceeds 80% of the actual maximum .
- The maximum amount of time during Pool BE Read Response Time exceeds 80% of the actual maximum.
- The maximum amount of time during Pool BE Write Response Time exceeds 80% of the actual maximum.

Use the **setdscfg** command to set or update the MDisk group attributes displayed by the **getdscfg** command:

- Example: `tpctool> setdscfg -property ManagedDiskGroupAvailIORate 99` This sets the ManagedDiskGroupAvailIORate value to 99% of the available capability for that managed-disk group.
- Output: No output is displayed for this command.

**Note:** The **getdscfg** and **setdscfg** commands can be used to display and modify other Tivoli Storage Productivity Center settings in addition to MDisk group settings.

- e. To create a storage resource group to report on VDisks by group, use the GUI or CLI.
- f. Create and run a performance monitor on the IBM System Storage SAN Volume Controller or IBM Storwize V7000 device.

**Note:** The performance monitor must be scheduled to run for at least the entire time period for which you want to capture report information.

## Installing Tivoli Common Reporting and running the storage tier reports

Once you have configured IBM Tivoli Storage Productivity Center to collect the appropriate storage tier data, install Tivoli Common Reporting and run the storage tier reports. The report results help you determine the actions you can take to improve your overall storage performance.

1. Install Tivoli Common Reporting. .  
For additional information about installing Tivoli Common Reporting, see the Tivoli Common Reporting Information Center and search for *installing*.
2. Log on to Tivoli Integrated Portal:
  - a. Open a browser and type `https://hostname:16311/ibm/console`.

**Tip:** 16311 is the default port number. If Tivoli Common Reporting and Tivoli Storage Productivity Center are installed on the same server, this port number might be different. For this type of configuration, use port 17310 to install Tivoli Common Reporting.

- b. Enter the user name and password.
3. In the navigation pane, expand **Reporting > Common Reporting**.
4. Select **TPC** and select a report from the list of available reports.
5. Select a storage virtualizer.
6. Specify the date range.
7. To run the report, click **Finish**.
8. View the report:
  - Specify report formats using the options in the upper right of the report.
  - Click **Page down** to see the next page of the report. Click **Bottom** to go to the end of the report.

## Running the storage tier reports

Storage tier reports can be run on schedules that are appropriate for the storage activity in your environment. If you have highly active storage and storage that continually exceeds the thresholds, you might have to run reports more frequently to identify potential data access problems. You can monitor active storage less frequently. You can easily report schedules and date ranges to match the activity in your storage environment.

## Interpreting the information displayed in storage tier reports

The columns in the Tivoli Storage Productivity Center storage tier reports provide information about the MDisk groups and VDIs in your storage environment. You can trace important performance data, such as read I/O capability, write I/O response time, and storage space capability. This information, combined with other data in the storage tier reports, can indicate areas of possible concern in your storage environment. The storage requirements for your specific environment can help you identify what report data is key for making storage tiering decisions:

- Is more I/O capability needed
- Is more storage capacity needed
- Do some applications require higher-tier storage on a seasonal basis?

Using the storage tier reports can help you answer these and other questions.

Use the information displayed in the storage tier reports provided by Tivoli Storage Productivity Center to help you improve your storage utilization. The reports are:

- MDisk Groups - Workload Activity Report
- MDisk Group Details Report
- MDisk Groups - VDisk Workload Activity Report
- MDisk Group - VDisk Workload Activity Details Report
- Storage Virtualizer - VDisk Workload Activity Report
- Storage Resource Group - VDisk Workload Activity Report
- Storage Resource Group - VDisk Workload Activity Details Report
- VDisk Details Report

## Restrictions when using storage tier reports

1. These reports cannot show performance data from releases prior to Tivoli Storage Productivity Center v4.2.2 since the Tivoli Storage Productivity Center performance monitor must be run for at least the number of days that you want to report on.

### Related tasks:

“Configure storage subsystems” on page 306

Use the Configure Devices wizard to set up storage subsystems for monitoring by Tivoli Storage Productivity Center. Storage subsystems that you configure in the wizard are automatically included in probe schedules and associated with alert conditions.

“Creating storage subsystem performance monitor” on page 377

Create a subsystem performance monitor to gather performance data for storage systems.

### Related reference:

“Create or Edit a Storage Resource Group” on page 504

Use the Create or Edit Storage Resource group panel to create or modify storage resource groups.

## MDisk Groups - Workload Activity Report

Use the MDisk Groups - Workload Activity Report to monitor the workload activity of managed disk (MDisk) groups.

The report identifies the least active and most active MDisk groups. It also provides information about the workload activity and capacity of MDisk groups. Based on the information that is provided in the report, you can analyze storage services in MDisk groups.

## Reporting periods

The name of the storage virtualizer, the date filter such as Last 7 days, and the reporting period are displayed at the top of the page.

If you select **All** as the search filter when you run a report, the system date is displayed as the start date for the report. Depending on the time zone, 1969 or 1970 is displayed in the year field of the start date.

## Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period of time. The default is 90 days. If you run a report after the data retention period has expired, the daily data for the MDisk group or the VDisk is unavailable.

## Read I/O rate capability

You must set the back-end pool properties of MDisk groups:

- To estimate the read I/O rate capability of an MDisk group
- To rank the workload activity of MDisk groups

If you do not set the back-end pool properties of an MDisk group:

- Not Set is displayed in the MDisk Group Read I/O Rate Capability column
- Not Available is displayed in the MDisk Group Read I/O Rate Availability column

If the read I/O rate capability of an MDisk group is less than the maximum read I/O rate of the MDisk group, negative values are displayed in the report.

## Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis is appended to the name. For example, if the name of an MDisk group is myverylongmanageddiskgroupname123 in a chart, it is truncated and displayed as myverylongmanageddiskgroupname...

Two charts are displayed:

### Most Active MDisk Groups chart

Most active MDisk groups are ranked using the following formula:  
 $\text{average}(\text{maximum\_read\_I/O\_rate} \div \text{read\_I/O\_rate\_capability})$

In this sample excerpt of report output for mdiskgroup1, the values that are required to rank most active MDisk groups are provided.

| MDisk Group Report Date | MDisk Group Read I/O Rate Capability (ops/s) | MDisk Group Max Read I/O Rate (ops/s) |
|-------------------------|----------------------------------------------|---------------------------------------|
| August 10 2011          | 552                                          | 243.79                                |
| August 11 2011          | 552                                          | 249.91                                |

Ranking is calculated as follows:

- For each report in an MDisk group, the value in the MDisk Group Max Read I/O Rate column is divided by the value in the MDisk Group Read I/O Rate Capability column:  
 $243.79 \div 552 = 0.44$   
 $249.91 \div 552 = 0.45$
- The result of the operation for each report is added:  
 $0.44 + 0.45 = 0.89$
- To get the average value, the total is divided by the number of days that are specified in the date range. For example, if the reporting period spans 2 days, the total is divided by 2.  
 $0.89 \div 2 = 0.45$
- The MDisk group with the highest average value is ranked 1. The MDisk group with the second highest average value is ranked 2, and so on.

### Least Active MDisk Groups chart

Least active MDisk groups are ranked by highest read I/O rate availability.

In this sample excerpt of report output for mdiskgroup2, the values that are required to rank least active MDisk Groups are provided.

| MDisk Group Report Date | MDisk Group Read I/O Rate Availability (ops/s) |
|-------------------------|------------------------------------------------|
| August 10 2011          | 198.00                                         |
| August 11 2011          | 199.00                                         |

Ranking is calculated as follows:

- For all of the reports for an MDisk group, the lowest value in the MDisk Group Read I/O Rate Availability column is used. In this sample, the lowest value for mdiskgroup2 is 198.00.
- The MDisk group with the highest read I/O availability rate is ranked 1. The MDisk group with the second highest read I/O availability rate is ranked 2, and so on.

The number of MDisk groups in a storage virtualizer determines how many MDisk groups are displayed in each chart, as shown in the following table.

| Number of MDisk groups in the storage virtualizer | Number of MDisk groups in the Most Active MDisk Groups chart | Number of MDisk groups in the Least Active MDisk Groups chart |
|---------------------------------------------------|--------------------------------------------------------------|---------------------------------------------------------------|
| ≥ 10                                              | 5                                                            | 5                                                             |
| 9                                                 | 5                                                            | 4                                                             |
| 8                                                 | 4                                                            | 4                                                             |
| 7                                                 | 4                                                            | 3                                                             |
| 6                                                 | 3                                                            | 3                                                             |
| 5                                                 | 3                                                            | 2                                                             |
| 4                                                 | 2                                                            | 2                                                             |
| 3                                                 | 2                                                            | 1                                                             |
| 2                                                 | 1                                                            | 1                                                             |
| 1                                                 | 1                                                            | 0                                                             |

**Tip:** Click the bar that represents an MDisk group in a chart to open the MDisk Group - VDisk Workload Activity Details Report. Alternatively, you can click the name of the MDisk group in the report.

### Threshold values

The names of the threshold values in the report that are used to monitor workload activity for MDisk groups are:

- Actual Size Threshold
- Read I/O Rate Threshold
- Max Read Response Time Threshold
- Max Write Response Time Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the value of the actual size threshold is changed during a reporting period from 50% to 60%, the name of the column is changed to MDisk Group Available Size < Actual Size Threshold of {50 - 60%}. The convention that is used to represent a threshold value is percent enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used:

$$(maximum\_response\_time\_value \times maximum\_response\_time\_threshold\_value) \div 100.$$

The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

## Report output

For each MDisk group, the following information is provided:

### **MDisk Group**

The name of the MDisk group.

### **MDisk Group Report Date**

The date that the data was collected for the report.

### **MDisk Group Capacity Status**

If the storage capacity that is allocated is less than the actual storage capacity, a status value of OK is displayed. If the storage capacity that is allocated is more than the actual storage capacity, a status value of Warning is displayed. The actual capacity of the MDisk group is calculated as follows:

$$(\text{mdisk\_group\_real\_capacity\_value} \div \text{mdisk\_group\_capacity\_value}) \times 100$$

### **MDisk Group Capacity (GB)**

The total storage capacity in GB of all of the disks in the MDisk group.

### **MDisk Group Real Capacity (GB)**

The total amount in GB of real storage capacity in the MDisk group.

### **MDisk Group Free Space (GB)**

The total amount of remaining capacity in GB that can be allocated to increase the size of an existing VDisk, or allocated to an additional VDisk.

### **MDisk Group Available Size < Actual Size Threshold of {%}**

If the amount of storage space that is available for allocation to VDIsks is less than the actual size threshold, Yes is displayed in the column. If the amount of storage space that is available for allocation to VDIsks is more than the actual size threshold, No is displayed in the column.

### **MDisk Group Read I/O Rate (ops/s)**

The read I/O rate of the MDisk group in operations per second.

### **MDisk Group Read I/O Rate Capability (ops/s)**

The read I/O capability rate of the MDisk group in operations per second.

**Tip:** To find out how the read I/O capability rate is calculated, see “Read I/O capability formula” on page 1129.

### **MDisk Group Max Read I/O Rate (ops/s)**

The maximum read I/O rate of the MDisk group in operations per second.

### **MDisk Group Read I/O Rate Availability (ops/s)**

The difference between the read I/O capability rate and the maximum read I/O rate in operations per second.

### **MDisk Group Read I/O Rate Availability < Read I/O Rate Capability Threshold of {%}**

The availability rate for read I/O operations is compared with the read I/O rate capability threshold. If the availability rate for read I/O operations is less than the read I/O rate capability threshold, Yes is displayed in the column. If the availability rate for read I/O operations is more than the read I/O rate capability threshold, No is displayed in the column.

### **MDisk Group Max Read Response Time (ms/op)**

The maximum read response time in milliseconds per operation.



**MDisk Group Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read response time threshold is exceeded over a 24 hour period.

**MDisk Group Average Read Response Time (ms/op)**

The average read response time in milliseconds per operation.

**MDisk Group Max Write I/O Rate (ops/s)**

The maximum write I/O rate in operations per second.

**MDisk Group Max Write Response Time (ms/op)**

The maximum write response time in milliseconds per operation.

**MDisk Group Max Write Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum write response time threshold is exceeded over a 24 hour period.

**MDisk Group Average Write Response Time (ms/op)**

The average write response time in milliseconds per operation.

**MDisk Group Average Write I/O Size (KB)**

The average size of write I/O operations in KB.

“Running storage tier reports” on page 1097

**MDisk Group Details Report**

Use the MDisk Group Details Report to monitor the workload activity of a managed disk (MDisk) group.

The report provides information about the workload activity and capacity of a selected MDisk group. Based on the information that is provided in the report, you can analyze storage services for an MDisk group.

**Reporting period**

The name of the storage virtualizer, the name of the MDisk group, the date filter such as Last 7 days, and the reporting period are displayed at the top of the page.

If you select **All** as the search filter when you run a report, the system date is displayed as the start date for the report. Depending on the time zone, 1969 or 1970 is displayed in the year field of the start date.

**Retention of report data**

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period of time. The default is 90 days. If you run a report after the data retention period has expired, the daily data for the MDisk group or the VDisk is unavailable.

**Read I/O rate capability**

If you do not set the back-end pool properties of an MDisk group:

- Not Set is displayed in the MDisk Group Read I/O Rate Capability column
- Not Available is displayed in the MDisk Group Read I/O Rate Availability column

If the read I/O rate capability of an MDisk group is less than the maximum read I/O rate of the MDisk group, negative values are displayed in the report.

## Threshold values

The names of the threshold values in the report that are used to monitor workload activity for MDisk groups are:

- Actual Size Threshold
- Read I/O Rate Threshold
- Max Read Response Time Threshold
- Max Write Response Time Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the value of the actual size threshold is changed during a reporting period from 50% to 60%, the name of the column is changed to MDisk Group Available Size < Actual Size Threshold of {50 - 60%}. The convention that is used to represent a threshold value is percent enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used:

$$(maximum\_response\_time\_value \times maximum\_response\_time\_threshold\_value) \div 100.$$

**Tip:** To monitor the workload activity of all MDisk groups in a storage virtualizer, run the MDisk Groups - Workload Activity report.

## Report output

The following information is provided about the MDisk group:

### **MDisk Group**

The name of the MDisk group.

### **MDisk Group Report Date**

The date that the data was collected for the report.

### **MDisk Group Capacity Status**

If the storage capacity that is allocated is less than the actual storage capacity, a status value of OK is displayed. If the storage capacity that is allocated is more than the actual storage capacity, a status value of Warning is displayed. The actual capacity of the MDisk group is calculated as follows:

$$(mdisk\_group\_real\_capacity\_value \div mdisk\_group\_capacity\_value) \times 100$$

### **MDisk Group Capacity (GB)**

The total storage capacity in GB of all of the disks in the MDisk group.

### **MDisk Group Real Capacity (GB)**

The total amount in GB of real storage capacity in the MDisk group.

### **MDisk Group Free Space (GB)**

The total amount of remaining capacity in GB that can be allocated to increase the size of an existing VDisk, or allocated to an additional VDisk.

### **MDisk Group Available Size < Actual Size Threshold of {%}**

If the amount of storage space that is available for allocation to VDIsks is less than the actual size threshold, Yes is displayed in the column. If the amount of storage space that is available for allocation to VDIsks is more than the actual size threshold, No is displayed in the column.

### **MDisk Group Read I/O Rate (ops/s)**

The read I/O rate of the MDisk group in operations per second.

**MDisk Group Read I/O Rate Capability (ops/s)**

The read I/O capability rate of the MDisk group in operations per second.

**Tip:** To find out how the read I/O capability rate is calculated, see “Read I/O capability formula” on page 1129.

**MDisk Group Max Read I/O Rate (ops/s)**

The maximum read I/O rate of the MDisk group in operations per second.

**MDisk Group Read I/O Rate Availability (ops/s)**

The difference between the read I/O capability rate and the maximum read I/O rate in operations per second.

**MDisk Group Read I/O Rate Availability < Read I/O Rate Capability Threshold of {%}**

The availability rate for read I/O operations is compared with the read I/O rate capability threshold. If the availability rate for read I/O operations is less than the read I/O rate capability threshold, Yes is displayed in the column. If the availability rate for read I/O operations is more than the read I/O rate capability threshold, No is displayed in the column.

**MDisk Group Max Read Response Time (ms/op)**

The maximum read response time in milliseconds per operation.

**MDisk Group Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read response time threshold is exceeded over a 24 hour period.

**MDisk Group Average Read Response Time (ms/op)**

The average read response time in milliseconds per operation.

**MDisk Group Max Write I/O Rate (ops/s)**

The maximum write I/O rate in operations per second.

**MDisk Group Max Write Response Time (ms/op)**

The maximum write response time in milliseconds per operation.

**MDisk Group Max Write Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum write response time threshold is exceeded over a 24 hour period.

**MDisk Group Average Write Response Time (ms/op)**

The average write response time in milliseconds per operation.

**MDisk Group Average Write I/O Size (KB)**

The average size of write I/O operations in KB.

“Running storage tier reports” on page 1097

Follow this procedure to run storage tier reports on managed disk (MDisk) groups, volumes (VDisks) volumes, and application storage usage and activity using Tivoli Common Reporting.

**MDisk Groups - VDisk Workload Activity Report**

Use the MDisk Groups - VDisk Workload Activity Report to monitor the performance of virtual disks (VDisks) in managed disk (MDisk) groups. The workload activity of each MDisk group is determined by calculating the average peak utilization of the VDIs in the MDisk group.

## Reporting period

The name of the storage virtualizer, the date filter such as Last 7 days, and the reporting period are displayed at the top of the page.

If you select **All** as the search filter when you run a report, the system date is displayed as the start date for the report. Depending on the time zone, 1969 or 1970 is displayed in the year field of the start date.

## Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period of time. The default is 90 days. If you run a report after the data retention period has expired, the daily data for the MDisk group or the VDisk is unavailable.

## VDisk migration

If a VDisk is migrated from one MDisk group to another MDisk group during a reporting period, daily VDisk reports are produced for both MDisk groups. Daily VDisk reports are produced for the source MDisk group from the start date of the report to the date the VDisk was migrated. Daily VDisk reports are also produced for the target MDisk group after the date the VDisk was migrated to the end date of the report. The ranking of most and least active VDIsks can be affected because the average peak utilization of the VDisk is calculated for the source MDisk group and the target MDisk group.

## Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis is appended to the name. For example, if the name of an MDisk group is myverylongmanageddiskgroupname123 in a chart, it is truncated and displayed as myverylongmanageddiskgroupname...

Two charts are displayed:

### Most Active MDisk Groups chart

Ranked by the highest average value of peak VDisk utilization in an MDisk group

### Least Active MDisk Groups chart

Ranked by the lowest average value of peak VDisk utilization in an MDisk group

The number of MDisk groups in the storage virtualizer determines how many MDisk groups are displayed in each chart, as shown in the following table.

| Number of MDisk groups in the storage virtualizer | Number of MDisk groups in the Most Active MDisk Groups chart | Number of MDisk groups in the Least Active MDisk Groups chart |
|---------------------------------------------------|--------------------------------------------------------------|---------------------------------------------------------------|
| ≥ 10                                              | 5                                                            | 5                                                             |
| 9                                                 | 5                                                            | 4                                                             |
| 8                                                 | 4                                                            | 4                                                             |
| 7                                                 | 4                                                            | 3                                                             |

| Number of MDisk groups in the storage virtualizer | Number of MDisk groups in the Most Active MDisk Groups chart | Number of MDisk groups in the Least Active MDisk Groups chart |
|---------------------------------------------------|--------------------------------------------------------------|---------------------------------------------------------------|
| 6                                                 | 3                                                            | 3                                                             |
| 5                                                 | 3                                                            | 2                                                             |
| 4                                                 | 2                                                            | 2                                                             |
| 3                                                 | 2                                                            | 1                                                             |
| 2                                                 | 1                                                            | 1                                                             |
| 1                                                 | 1                                                            | 0                                                             |

**Tip:** Click the bar that represents an MDisk group in a chart to open the MDisk Group - VDisk Workload Activity Details Report. Alternatively, click the name of the MDisk group in the report. Click the name of the VDisk in the report to open the VDisk Details Report.

### Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDIs are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is percent enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used:  $(\text{maximum\_response\_time\_value} \times \text{maximum\_response\_time\_threshold\_value}) / 100$ . The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

### Report output

For each VDisk in an MDisk group, the following information is provided:

**MDisk Group**

The name of the MDisk group.

**VDisk** The name of the VDisk.

**VDisk Report Date**

The date that the data was collected for the report.

**VDisk Capacity (GB)**

The total storage capacity of the VDisk in GB.

**VDisk Thin Provisioning Status**

The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is  $\geq$  the specified warning threshold value, a status value of Warning is displayed. If the used capacity of the VDisk in GB is  $<$  the specified warning threshold value, a status value of OK is displayed. If thin provisioning is not used, a status value of Not Applicable is displayed.

**VDisk Peak Utilization (%)**

The highest value of VDisk utilization over a 24-hour period.

**VDisk Read I/O rate (ops/s)**

The read I/O rate of the VDisk in operations per second.

**VDisk Max Read I/O rate (ops/s)**

The maximum read I/O rate of the VDisk in operations per second.

**VDisk Average Read Response Time (ms/op)**

The average read response time in milliseconds per operation.

**VDisk Max Read Response Time (ms/op)**

The maximum read response time in milliseconds per operation.

**VDisk Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

**VDisk Write I/O Rate (ops/s)**

The write I/O rate in operations per second.

**VDisk Average Write Response Time (ms/op)**

The average write response time in milliseconds per operation.

**VDisk Max Write Response Time (ms/op)**

The maximum write response time in milliseconds per operation.

**VDisk Max Write Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

**VDisk Max Read Density Rate (ops/s per GB)**

The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk has allocated.

**VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)**

The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDIs in the MDisk group.

**VDisk Capacity Share of Total VDisk Capacity (%)**

The capacity of the VDisk as a percentage of the total capacity of all the VDIs in the MDisk group.

**VDisk Average Read Cache Hits (%)**

The average read cache hits in percent.

**VDisk Max Read Cache Hits (%)**

The maximum read cache hits in percent.

**VDisk Max Read Cache Hits Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

**VDisk Max Read I/O Rate Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

**VDisk Max Write I/O Rate Threshold {%} Exceeded of By (hr.:min.)**

The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

**VDisk Max Write Cache Delay I/O Rate (ops/s)**

The maximum write cache delay I/O rate of the VDisk in operations per second.

“Running storage tier reports” on page 1097

**MDisk Group - VDisk Workload Activity Details Report**

Use the MDisk Group - VDisk Workload Activity Details Report to monitor the workload activity of virtual disks (VDisks) in a managed disk (MDisk) group. The workload activity of VDisks in an MDisk group is determined by calculating the average peak utilization of each VDisk.

**Reporting period**

The name of the storage virtualizer, the name of the MDisk group, the date filter such as Last 7 days, and the reporting period are displayed at the top of the page.

If you select **All** as the search filter when you run a report, the system date is displayed as the start date for the report. Depending on the time zone, 1969 or 1970 is displayed in the year field of the start date.

**Retention of report data**

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period of time. The default is 90 days. If you run a report after the data retention period has expired, the daily data for the MDisk group or the VDisk is unavailable.

**VDisk migration**

If a VDisk is migrated from one MDisk group to another MDisk group during a reporting period, daily VDisk reports are produced for both MDisk groups. Daily VDisk reports are produced for the source MDisk group from the start date of the report to the date the VDisk was migrated. Daily VDisk reports are also produced for the target MDisk group after the date the VDisk was migrated to the end date of the report. The ranking of most and least active VDisks can be affected because the average peak utilization of the VDisk is calculated for the source MDisk group and the target MDisk group.

**Charts**

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis is appended to the name. For example, if the name of an MDisk group is myverylongmanageddiskgroupname123 in a chart, it is truncated and displayed as myverylongmanageddiskgroupname...

Two charts are displayed:

### Most Active VDIs chart

Ranked by the highest average value of peak VDisk utilization in the MDisk group

### Least Active VDIs chart

Ranked by the lowest average value of peak VDisk utilization in the MDisk group

The number of VDIs in an MDisk group determines how many VDIs are displayed in each chart, as shown in the following table.

| Number of VDIs in the MDisk group | Number of VDIs in the Most Active VDIs chart | Number of VDIs in the Least Active VDIs chart |
|-----------------------------------|----------------------------------------------|-----------------------------------------------|
| ≥ 10                              | 5                                            | 5                                             |
| 9                                 | 5                                            | 4                                             |
| 8                                 | 4                                            | 4                                             |
| 7                                 | 4                                            | 3                                             |
| 6                                 | 3                                            | 3                                             |
| 5                                 | 3                                            | 2                                             |
| 4                                 | 2                                            | 2                                             |
| 3                                 | 2                                            | 1                                             |
| 2                                 | 1                                            | 1                                             |
| 1                                 | 1                                            | 0                                             |

**Tip:** Click the bar that represents a VDisk in a chart to open the VDisk Details Report. Alternatively, click the name of the VDisk in the report. Click the name of the MDisk group in the report to open the MDisk Group Details Report.

### Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDIs are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is percent enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used:  $(\text{maximum\_response\_time\_value} \times \text{maximum\_response\_time\_threshold\_value})/100$ . The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.



## Report output

For each VDisk in an MDisk group, the following information is provided:

### **MDisk Group**

The name of the MDisk group.

**VDisk** The name of the VDisk.

### **VDisk Report Date**

The date that the data was collected for the report.

### **VDisk Capacity (GB)**

The total storage capacity of the VDisk in GB.

### **VDisk Thin Provisioning Status**

The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is  $\geq$  the specified warning threshold value, a status value of `Warning` is displayed. If the used capacity of the VDisk in GB is  $<$  the specified warning threshold value, a status value of `OK` is displayed. If thin provisioning is not used, a status value of `Not Applicable` is displayed.

### **VDisk Peak Utilization (%)**

The highest value of VDisk utilization over a 24-hour period.

### **VDisk Read I/O rate (ops/s)**

The read I/O rate of the VDisk in operations per second.

### **VDisk Max Read I/O rate (ops/s)**

The maximum read I/O rate of the VDisk in operations per second.

### **VDisk Average Read Response Time (ms/op)**

The average read response time in milliseconds per operation.

### **VDisk Max Read Response Time (ms/op)**

The maximum read response time in milliseconds per operation.

### **VDisk Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

### **VDisk Write I/O Rate (ops/s)**

The write I/O rate in operations per second.

### **VDisk Average Write Response Time (ms/op)**

The average write response time in milliseconds per operation.

### **VDisk Max Write Response Time (ms/op)**

The maximum write response time in milliseconds per operation.

### **VDisk Max Write Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

### **VDisk Max Read Density Rate (ops/s per GB)**

The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk has allocated.

**VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)**

The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDIs in the MDisk group.

**VDisk Capacity Share of Total VDisk Capacity (%)**

The capacity of the VDisk as a percentage of the total capacity of all the VDIs in the MDisk group.

**VDisk Average Read Cache Hits (%)**

The average read cache hits in percent.

**VDisk Max Read Cache Hits (%)**

The maximum read cache hits in percent.

**VDisk Max Read Cache Hits Time Threshold of { } Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

**VDisk Max Read I/O Rate Threshold of { } Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

**VDisk Max Write I/O Rate Threshold { } Exceeded of By (hr.:min.)**

The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

**VDisk Max Write Cache Delay I/O Rate (ops/s)**

The maximum write cache delay I/O rate of the VDisk in operations per second.

“Running storage tier reports” on page 1097

Follow this procedure to run storage tier reports on managed disk (MDisk) groups, volumes (VDisks) volumes, and application storage usage and activity using Tivoli Common Reporting.

**Storage Virtualizer - VDisk Workload Activity Report**

Use the Storage Virtualizer - VDisk Workload Activity Report to monitor the workload activity of virtual disks (VDisks) in a storage virtualizer. The workload activity of VDIs in a storage virtualizer is determined by calculating the average peak utilization of each VDisk.

**Reporting period**

The name of the storage virtualizer, the date filter such as Last 7 days, and the reporting period are displayed at the top of the page.

If you select **All** as the search filter when you run a report, the system date is displayed as the start date for the report. Depending on the time zone, 1969 or 1970 is displayed in the year field of the start date.

**Retention of report data**

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period of time. The default is 90 days. If you run a report after the data retention period has expired, the daily data for the MDisk group or the VDisk is unavailable.

## VDisk migration

If a VDisk is migrated from one MDisk group to another MDisk group during a reporting period, daily VDisk reports are produced for both MDisk groups. Daily VDisk reports are produced for the source MDisk group from the start date of the report to the date the VDisk was migrated. Daily VDisk reports are also produced for the target MDisk group after the date the VDisk was migrated to the end date of the report. The ranking of most and least active VDIs can be affected because the average peak utilization of the VDisk is calculated for the source MDisk group and the target MDisk group.

## Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis is appended to the name. For example, if the name of an MDisk group is myverylongmanageddiskgroupname123 in a chart, it is truncated and displayed as myverylongmanageddiskgroupname...

Two charts are displayed:

### Most Active VDIs chart

Ranked by the highest average value of peak VDisk utilization in the storage virtualizer

### Least Active VDIs chart

Ranked by the lowest average value of peak VDisk utilization in the storage virtualizer

The number of VDIs in a storage virtualizer determines how many VDIs are displayed in each chart, as shown in the following table.

| Number of VDIs in the storage virtualizer | Number of VDIs in the Most Active VDIs chart | Number of VDIs in the Least Active VDIs chart |
|-------------------------------------------|----------------------------------------------|-----------------------------------------------|
| ≥ 10                                      | 5                                            | 5                                             |
| 9                                         | 5                                            | 4                                             |
| 8                                         | 4                                            | 4                                             |
| 7                                         | 4                                            | 3                                             |
| 6                                         | 3                                            | 3                                             |
| 5                                         | 3                                            | 2                                             |
| 4                                         | 2                                            | 2                                             |
| 3                                         | 2                                            | 1                                             |
| 2                                         | 1                                            | 1                                             |
| 1                                         | 1                                            | 0                                             |

**Tip:** Click the bar that represents a VDisk in a chart to open the VDisk Details Report. Alternatively, click the name of the VDisk in the report. Click the name of the MDisk group in the report to open the MDisk Group Details Report.

## Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDIs are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is percent enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used:  $(\text{maximum\_response\_time\_value} \times \text{maximum\_response\_time\_threshold\_value})/100$ . The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

## Report output

For each VDisk in a storage virtualizer, the following information is displayed:

### **MDisk Group**

The name of the MDisk group.

**VDisk** The name of the VDisk.

### **VDisk Report Date**

The date that the data was collected for the report.

### **VDisk Capacity (GB)**

The total storage capacity of the VDisk in GB.

### **VDisk Thin Provisioning Status**

The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is  $\geq$  the specified warning threshold value, a status value of Warning is displayed. If the used capacity of the VDisk in GB is  $<$  the specified warning threshold value, a status value of OK is displayed. If thin provisioning is not used, a status value of Not Applicable is displayed.

### **VDisk Peak Utilization (%)**

The highest value of VDisk utilization over a 24-hour period.

### **VDisk Read I/O rate (ops/s)**

The read I/O rate of the VDisk in operations per second.

### **VDisk Max Read I/O rate (ops/s)**

The maximum read I/O rate of the VDisk in operations per second.

### **VDisk Average Read Response Time (ms/op)**

The average read response time in milliseconds per operation.

### **VDisk Max Read Response Time (ms/op)**

The maximum read response time in milliseconds per operation.

**VDisk Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

**VDisk Write I/O Rate (ops/s)**

The write I/O rate in operations per second.

**VDisk Average Write Response Time (ms/op)**

The average write response time in milliseconds per operation.

**VDisk Max Write Response Time (ms/op)**

The maximum write response time in milliseconds per operation.

**VDisk Max Write Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

**VDisk Max Read Density Rate (ops/s per GB)**

The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk has allocated.

**VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)**

The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDIs in the MDisk group.

**VDisk Capacity Share of Total VDisk Capacity (%)**

The capacity of the VDisk as a percentage of the total capacity of all the VDIs in the MDisk group.

**VDisk Average Read Cache Hits (%)**

The average read cache hits in percent.

**VDisk Max Read Cache Hits (%)**

The maximum read cache hits in percent.

**VDisk Max Read Cache Hits Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

**VDisk Max Read I/O Rate Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

**VDisk Max Write I/O Rate Threshold {%} Exceeded of By (hr.:min.)**

The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

**VDisk Max Write Cache Delay I/O Rate (ops/s)**

The maximum write cache delay I/O rate of the VDisk in operations per second.

“Running storage tier reports” on page 1097

Follow this procedure to run storage tier reports on managed disk (MDisk) groups, volumes (VDisks) volumes, and application storage usage and activity using Tivoli Common Reporting.

## **Storage Resource Group - VDisk Workload Activity Report**

Use the Storage Resource Group - VDisk Workload Activity Report to monitor the workload activity of virtual disks (VDisks) in storage resource groups. The workload activity of each storage resource group is determined by calculating the average peak utilization of the VDIs in the storage resource group.

## Reporting period

The name of the storage virtualizer, the date filter such as Last 7 days, and the reporting period are displayed at the top of the page.

If you select **All** as the search filter when you run a report, the system date is displayed as the start date for the report. Depending on the time zone, 1969 or 1970 is displayed in the year field of the start date.

## Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period of time. The default is 90 days. If you run a report after the data retention period has expired, the daily data for the MDisk group or the VDisk is unavailable.

## VDisk migration

If a VDisk is migrated from one MDisk group to another MDisk group during a reporting period, daily VDisk reports are produced for both MDisk groups. Daily VDisk reports are produced for the source MDisk group from the start date of the report to the date the VDisk was migrated. Daily VDisk reports are also produced for the target MDisk group after the date the VDisk was migrated to the end date of the report. The ranking of most and least active VDIsks can be affected because the average peak utilization of the VDisk is calculated for the source MDisk group and the target MDisk group.

## Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis is appended to the name. For example, if the name of an MDisk group is myverylongmanageddiskgroupname123 in a chart, it is truncated and displayed as myverylongmanageddiskgroupname...

Two charts are displayed:

### Most Active Storage Resource Groups chart

Ranked by the highest average value of peak VDisk utilization in storage resource groups

### Least Active Storage Resource Groups chart

Ranked by the lowest average value of peak VDisk utilization in storage resource groups

The number of storage resource groups in a storage virtualizer determines how many storage resource groups are displayed in each chart, as shown in the table.

| Number of storage resource groups in the storage virtualizer | Number of storage resource groups in the Most Active Storage Resource Groups chart | Number of storage resource groups in the Least Active Storage Resource Groups chart |
|--------------------------------------------------------------|------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|
| ≥ 10                                                         | 5                                                                                  | 5                                                                                   |
| 9                                                            | 5                                                                                  | 4                                                                                   |
| 8                                                            | 4                                                                                  | 4                                                                                   |
| 7                                                            | 4                                                                                  | 3                                                                                   |

| Number of storage resource groups in the storage virtualizer | Number of storage resource groups in the Most Active Storage Resource Groups chart | Number of storage resource groups in the Least Active Storage Resource Groups chart |
|--------------------------------------------------------------|------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|
| 6                                                            | 3                                                                                  | 3                                                                                   |
| 5                                                            | 3                                                                                  | 2                                                                                   |
| 4                                                            | 2                                                                                  | 2                                                                                   |
| 3                                                            | 2                                                                                  | 1                                                                                   |
| 2                                                            | 1                                                                                  | 1                                                                                   |
| 1                                                            | 1                                                                                  | 0                                                                                   |

**Tip:** Click the bar that represents a storage resource group to open the Storage Resource Group: VDisk Workload Activity Details Report. Alternatively, click the name of the storage resource group in the report. Click the name of the VDisk in the report to open the VDisk Details Report. Click the name of the MDisk group in the report to open the MDisk Group Details Report.

### Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDIs are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is percent enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used:  $(\text{maximum\_response\_time\_value} \times \text{maximum\_response\_time\_threshold\_value})/100$ . The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

### Report output

For each storage resource group in a storage virtualizer, the following information is provided about VDIs:

#### Storage Resource Group

The name of the storage resource group.

#### MDisk Group

The name of the MDisk group.

**VDisk** The name of the VDisk.

**VDisk Report Date**

The date that the data was collected for the report.

**VDisk Capacity (GB)**

The total storage capacity of the VDisk in GB.

**VDisk Thin Provisioning Status**

The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is  $\geq$  the specified warning threshold value, a status value of `Warning` is displayed. If the used capacity of the VDisk in GB is  $<$  the specified warning threshold value, a status value of `OK` is displayed. If thin provisioning is not used, a status value of `Not Applicable` is displayed.

**VDisk Peak Utilization (%)**

The highest value of VDisk utilization over a 24-hour period.

**VDisk Read I/O rate (ops/s)**

The read I/O rate of the VDisk in operations per second.

**VDisk Max Read I/O rate (ops/s)**

The maximum read I/O rate of the VDisk in operations per second.

**VDisk Average Read Response Time (ms/op)**

The average read response time in milliseconds per operation.

**VDisk Max Read Response Time (ms/op)**

The maximum read response time in milliseconds per operation.

**VDisk Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

**VDisk Write I/O Rate (ops/s)**

The write I/O rate in operations per second.

**VDisk Average Write Response Time (ms/op)**

The average write response time in milliseconds per operation.

**VDisk Max Write Response Time (ms/op)**

The maximum write response time in milliseconds per operation.

**VDisk Max Write Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

**VDisk Max Read Density Rate (ops/s per GB)**

The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk has allocated.

**VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)**

The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDIs in the MDisk group.

**VDisk Capacity Share of Total VDisk Capacity (%)**

The capacity of the VDisk as a percentage of the total capacity of all the VDIs in the MDisk group.

**VDisk Average Read Cache Hits (%)**

The average read cache hits in percent.



**VDisk Max Read Cache Hits (%)**

The maximum read cache hits in percent.

**VDisk Max Read Cache Hits Time Threshold of { } Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

**VDisk Max Read I/O Rate Threshold of { } Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

**VDisk Max Write I/O Rate Threshold { } Exceeded of By (hr.:min.)**

The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

**VDisk Max Write Cache Delay I/O Rate (ops/s)**

The maximum write cache delay I/O rate of the VDisk in operations per second.

“Running storage tier reports” on page 1097

Follow this procedure to run storage tier reports on managed disk (MDisk) groups, volumes (VDisks) volumes, and application storage usage and activity using Tivoli Common Reporting.

**Storage Resource Group - VDisk Workload Activity Details Report**

Use the Storage Resource Group - VDisk Workload Activity Details Report to monitor the workload activity of virtual disks (VDisks) in a storage resource group. The workload activity of VDisks in a storage resource group, such as a database application, is determined by calculating the average peak utilization of each VDisk.

**Reporting period**

The name of the storage virtualizer, the name of the storage resource group, the date filter such as Last 7 days, and the reporting period are displayed at the top of the page.

If you select **All** as the search filter when you run a report, the system date is displayed as the start date for the report. Depending on the time zone, 1969 or 1970 is displayed in the year field of the start date.

**Retention of report data**

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period of time. The default is 90 days. If you run a report after the data retention period has expired, the daily data for the MDisk group or the VDisk is unavailable.

**VDisk migration**

If a VDisk is migrated from one MDisk group to another MDisk group during a reporting period, daily VDisk reports are produced for both MDisk groups. Daily VDisk reports are produced for the source MDisk group from the start date of the report to the date the VDisk was migrated. Daily VDisk reports are also produced for the target MDisk group after the date the VDisk was migrated to the end date of the report. The ranking of most and least active VDisks can be affected because the average peak utilization of the VDisk is calculated for the source MDisk group and the target MDisk group.

## Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis is appended to the name. For example, if the name of an MDisk group is myverylongmanageddiskgroupname123 in a chart, it is truncated and displayed as myverylongmanageddiskgroupname...

Two charts are displayed:

### Most Active VDIsks chart

Ranked by the highest average value of peak VDisk utilization in the storage resource group

### Least Active VDIsks chart

Ranked by the lowest average value of peak VDisk utilization in the storage resource group

The number of VDIsks that a storage resource group uses determines how many VDIsks are displayed in each chart, as shown in the table.

| Number of VDIsks that the storage resource group uses | Number of VDIsks in the Most Active VDIsks chart | Number of VDIsks in the Least Active VDIsks chart |
|-------------------------------------------------------|--------------------------------------------------|---------------------------------------------------|
| ≥ 10                                                  | 5                                                | 5                                                 |
| 9                                                     | 5                                                | 4                                                 |
| 8                                                     | 4                                                | 4                                                 |
| 7                                                     | 4                                                | 3                                                 |
| 6                                                     | 3                                                | 3                                                 |
| 5                                                     | 3                                                | 2                                                 |
| 4                                                     | 2                                                | 2                                                 |
| 3                                                     | 2                                                | 1                                                 |
| 2                                                     | 1                                                | 1                                                 |
| 1                                                     | 1                                                | 0                                                 |

**Tip:** Click the bar that represents a VDisk in a chart to open the VDisk Details Report. Alternatively, click the name of the VDisk in the report. Click the name of the MDisk group in the report to open the MDisk Group Details Report. To monitor the workload activity of VDIsks for all storage resource groups in a storage virtualizer, run the Storage Resource Group - VDisk Workload Activity Report.

## Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDIsks are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is percent enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used:  $(\text{maximum\_response\_time\_value} \times \text{maximum\_response\_time\_threshold\_value})/100$ . The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

## Report output

The following information is provided about VDIs in the storage resource group:

### **MDisk Group**

The name of the MDisk group.

**VDisk** The name of the VDisk.

### **VDisk Report Date**

The date that the data was collected for the report.

### **VDisk Capacity (GB)**

The total storage capacity of the VDisk in GB.

### **VDisk Thin Provisioning Status**

The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is  $\geq$  the specified warning threshold value, a status value of Warning is displayed. If the used capacity of the VDisk in GB is  $<$  the specified warning threshold value, a status value of OK is displayed. If thin provisioning is not used, a status value of Not Applicable is displayed.

### **VDisk Peak Utilization (%)**

The highest value of VDisk utilization over a 24-hour period.

### **VDisk Read I/O rate (ops/s)**

The read I/O rate of the VDisk in operations per second.

### **VDisk Max Read I/O rate (ops/s)**

The maximum read I/O rate of the VDisk in operations per second.

### **VDisk Average Read Response Time (ms/op)**

The average read response time in milliseconds per operation.

### **VDisk Max Read Response Time (ms/op)**

The maximum read response time in milliseconds per operation.

### **VDisk Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

### **VDisk Write I/O Rate (ops/s)**

The write I/O rate in operations per second.

**VDisk Average Write Response Time (ms/op)**

The average write response time in milliseconds per operation.

**VDisk Max Write Response Time (ms/op)**

The maximum write response time in milliseconds per operation.

**VDisk Max Write Response Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

**VDisk Max Read Density Rate (ops/s per GB)**

The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk has allocated.

**VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)**

The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDIs in the MDisk group.

**VDisk Capacity Share of Total VDisk Capacity (%)**

The capacity of the VDisk as a percentage of the total capacity of all the VDIs in the MDisk group.

**VDisk Average Read Cache Hits (%)**

The average read cache hits in percent.

**VDisk Max Read Cache Hits (%)**

The maximum read cache hits in percent.

**VDisk Max Read Cache Hits Time Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

**VDisk Max Read I/O Rate Threshold of {%} Exceeded By (hr.:min.)**

The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

**VDisk Max Write I/O Rate Threshold {%} Exceeded of By (hr.:min.)**

The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

**VDisk Max Write Cache Delay I/O Rate (ops/s)**

The maximum write cache delay I/O rate of the VDisk in operations per second.

“Running storage tier reports” on page 1097

Follow this procedure to run storage tier reports on managed disk (MDisk) groups, volumes (VDisks) volumes, and application storage usage and activity using Tivoli Common Reporting.

**VDisk Details Report**

Use the VDisk Details Report to monitor the workload activity of a virtual disk (VDisk) in a managed disk group.

Based on the information that the report provides, you can analyze the performance of a VDisk in a managed disk (MDisk) group.

**Reporting period**

If you select **All** as the search filter when you run a report, the system date is displayed as the start date for the report. Depending on the time zone, 1969 or 1970 is displayed in the year field of the start date.

## Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period of time. The default is 90 days. If you run a report after the data retention period has expired, the daily data for the MDisk group or the VDisk is unavailable.

## Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDIs are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is percent enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used:  $(\text{maximum\_response\_time\_value} \times \text{maximum\_response\_time\_threshold\_value})/100$ . The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

## Report output

The following information is displayed about the VDisk:

### **MDisk Group**

The name of the MDisk group.

**VDisk** The name of the VDisk.

### **VDisk Report Date**

The date that the data was collected for the report.

### **VDisk Capacity (GB)**

The total storage capacity of the VDisk in GB.

### **VDisk Thin Provisioning Status**

The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is  $\geq$  the specified warning threshold value, a status value of Warning is displayed. If the used capacity of the VDisk in GB is  $<$  the specified warning threshold value, a status value of OK is displayed. If thin provisioning is not used, a status value of Not Applicable is displayed.

### **VDisk Peak Utilization (%)**

The highest value of VDisk utilization over a 24-hour period.

- VDisk Read I/O rate (ops/s)**  
The read I/O rate of the VDisk in operations per second.
- VDisk Max Read I/O rate (ops/s)**  
The maximum read I/O rate of the VDisk in operations per second.
- VDisk Average Read Response Time (ms/op)**  
The average read response time in milliseconds per operation.
- VDisk Max Read Response Time (ms/op)**  
The maximum read response time in milliseconds per operation.
- VDisk Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)**  
The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.
- VDisk Write I/O Rate (ops/s)**  
The write I/O rate in operations per second.
- VDisk Average Write Response Time (ms/op)**  
The average write response time in milliseconds per operation.
- VDisk Max Write Response Time (ms/op)**  
The maximum write response time in milliseconds per operation.
- VDisk Max Write Response Time Threshold of {%} Exceeded By (hr.:min.)**  
The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.
- VDisk Max Read Density Rate (ops/s per GB)**  
The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk has allocated.
- VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)**  
The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDIsks in the MDisk group.
- VDisk Capacity Share of Total VDisk Capacity (%)**  
The capacity of the VDisk as a percentage of the total capacity of all the VDIsks in the MDisk group.
- VDisk Average Read Cache Hits (%)**  
The average read cache hits in percent.
- VDisk Max Read Cache Hits (%)**  
The maximum read cache hits in percent.
- VDisk Max Read Cache Hits Time Threshold of {%} Exceeded By (hr.:min.)**  
The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.
- VDisk Max Read I/O Rate Threshold of {%} Exceeded By (hr.:min.)**  
The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.
- VDisk Max Write I/O Rate Threshold {%} Exceeded of By (hr.:min.)**  
The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.
- VDisk Max Write Cache Delay I/O Rate (ops/s)**  
The maximum write cache delay I/O rate of the VDisk in operations per second.

“Running storage tier reports” on page 1097

Follow this procedure to run storage tier reports on managed disk (MDisk) groups, volumes (VDisks) volumes, and application storage usage and activity using Tivoli Common Reporting.

### Read I/O capability formula

The read I/O capability formula is used by IBM Tivoli Storage Productivity Center to estimate the workload capacity of managed disk (MDisk) groups.

The estimated workload capacity is compared with the actual workload capacity to assess the most active and least active MDisk groups in a storage subsystem. To monitor the workload activity of MDisk groups, you can run customized reports such as the MDisk Groups - Workload Activity Report. This formula is used for MDisk groups in the following storage systems:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified

The formula that is used by IBM Tivoli Storage Productivity Center to estimate the read I/O capability rate is:

$$\frac{[\text{Estimated Read I/O Workload} * \text{Type of Disk I/O Capability} * \text{Number of Disks}]}{[\text{Estimated Read I/O Workload} * (1 - \text{Cache Hit Ratio}) + \text{Estimated Write I/O Workload} * \text{RAID Type Weighted I/O Ratio}]}$$

There are two types of value that are used by the formula:

- Fixed values provided by IBM Tivoli Storage Productivity Center
- Variable values that you must specify for each MDisk group

The following fixed values are used by the formula:

#### Estimated Read I/O Workload

The estimated read I/O workload of an MDisk group. The estimate is set to a fixed value of 0.7 (70%).

#### Estimated Write I/O Workload

The estimated write I/O workload of an MDisk group. The estimate is set to a fixed value of 0.3 (30%).

You specify the following variable values for each MDisk group:

#### Type of Disk I/O Capability

When you specify the type of disk that the back-end pool uses:

- You can accept the default I/O capability per second of the disk
- You can enter a new value for the I/O capability per second of the disk

For example, the default value for the I/O capability rate per second of fiber types of disk with 15,000 revolutions per minute (rpm) is 150 I/O operations per second. To change the I/O capability for a type of disk, you issue the **setbackenddisktype** command and you enter a new value for the **-iops** parameter.

#### Number of Disks

You specify the number of physical disks that the back-end pool uses.

#### Cache Hit Ratio

When you specify the type of storage system that the back-end pool uses:

- You can accept the default cache hit ratio for the storage system
- You can enter a new value for the cache hit ratio of the storage system

For example, the default cache hit ratio of the XIV system is 0.7 (70%). To change the cache hit ratio for a type of storage system, you issue the **setbackendtype** command and you enter a new value for the **-cachehit** parameter.

#### RAID Type Weighted I/O Ratio

When you specify the type of RAID that the back-end pool uses:

- You can accept the default value for the weighted I/O ratio of the type of RAID
- You can enter a new value for the weighted I/O ratio of the type of RAID

For example the weighted I/O ratio for RAID5 types of RAID is 4:1. To change the weighted I/O ratio for a type of RAID, you issue the **setbackendraidtype** command and you enter a new value for the **-weightedio** parameter.

**Tip:** To set the values of back-end pools for an MDisk group, access the MDisk Group Details page through the **Disk Manager > MDisk Groups** node of the navigation tree. You can also use the Tivoli Storage Productivity Center CLI called **tcptool** to set the values for back-end pools. See the **setarray** command. To change the variable values that are used by the Read I/O capability formula, you must use the Tivoli Storage Productivity Center CLI.

#### Example

The following values have been set for a back-end pool. The type of storage system is the XIV system. The type of disk is F15 (Fiber - 15 000 rpm), and the type of RAID is X.

| Name of variable             | Value     |
|------------------------------|-----------|
| Cache Hit Ratio              | 0.5 (50%) |
| Type of Disk I/O Capability  | 150       |
| RAID Type Weighted I/O Ratio | 4         |
| Number of Disks              | 10        |

Based on the values that are provided, the read I/O capability per second is:

$$[0.7 * 150 * 10] / [0.7 * (1 - 0.5) + 0.3 * 4]$$

The read I/O capability per second rate is 677. (The result of the calculation is rounded to the nearest whole number.)

---

## Tivoli Storage Productivity Center Monitoring Agent reports

After you have installed and configured the Tivoli Storage Productivity Center Monitoring Agent, you can view reports and information from Tivoli Storage Productivity Center.

There are four predefined workspaces which are shipped with Tivoli Storage Productivity Center:



## TPC Server Overview

Provides an over view of the key parameters and infrastructure metrics about the Data server and Device server managed entities.

## Managed SAN Overview

Provides an overview about several metrics concerning the managed SAN environment.

## Managed SAN Health

Monitors the health and availability of the SAN attached devices managed by Tivoli Storage Productivity Center.

## TPC Data Collection Overview

Displays the status of the latest probe jobs and performance monitors.

## Managed SAN Overview workspace

The Managed SAN Overview workspace provides an overview about several metrics about the managed SAN environment. You can access this workspace by clicking **Managed SAN Overview** in the navigation tree of the Navigator window.

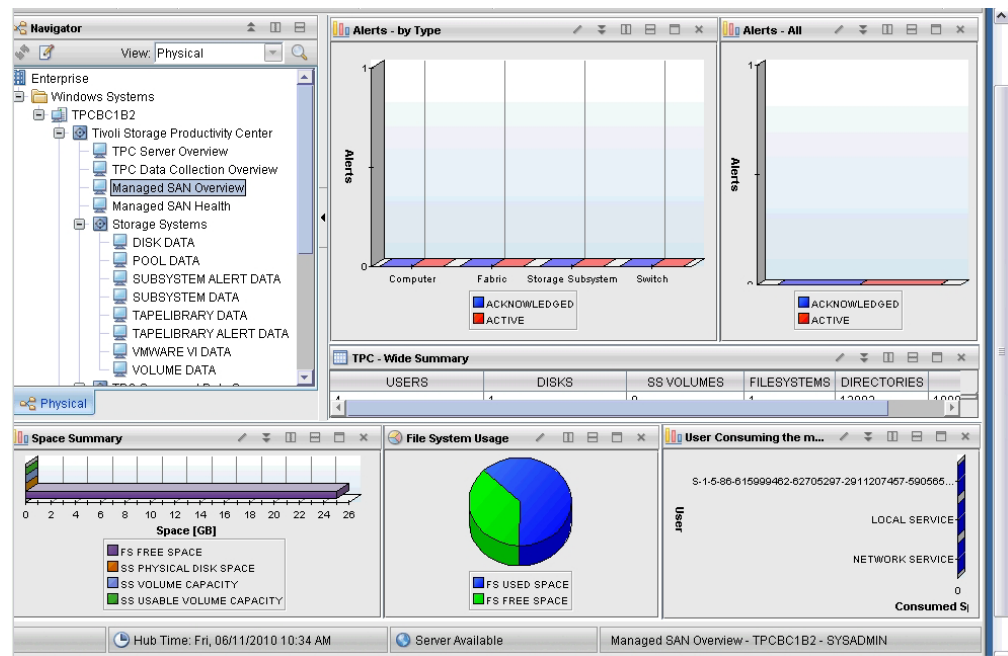


Figure 61. Managed SAN overview report

You can see the following views:

### Alerts - by Type

Displays a bar chart showing the number of all computers, switches, storage systems, and fabric alerts including their current state (acknowledged or active).

### Alerts - All

Displays a bar chart showing the number of all alerts of all devices including their current state (acknowledged or active).

### TPC-Wide Summary

Displays a table showing the total amount of users, disks, storage system volumes, file systems, files, and folders.

## Space Summary

Displays a bar chart showing the total amount of file system capacity, physical disk space of all storage systems, total amount of the available volume capacity, and total amount of usable volume capacity in the managed storage environment.

## User Consuming the most Space

Displays a chart with the total amount of space allocated by single users. Only the ten users using the most space is shown.

## File System Usage

Displays a pie chart showing the total amount of free and used file system space in the managed storage environment.

## Managed SAN Health workspace

The Managed SAN Health workspace provides an overview about the objects in the managed SAN environment. You can access this workspace by clicking **Managed SAN Health** in the navigation tree of the Navigator window.

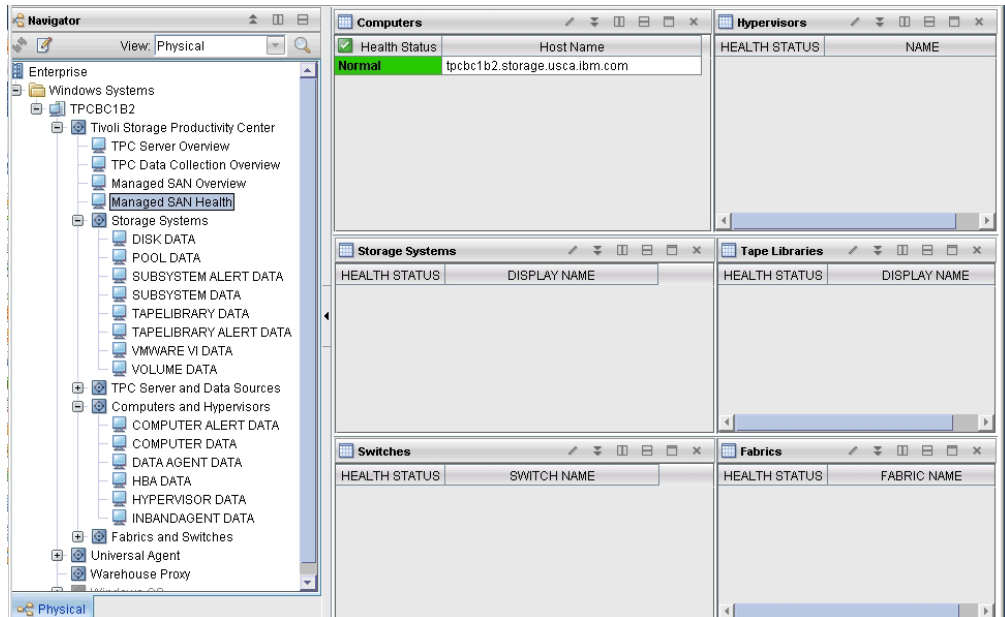


Figure 62. Managed SAN Health report

This workspace shows the following information:

### Computers

Displays a list of all managed computers including their health status.

### Hypervisors

Displays a list of all managed VMware hypervisors including their health status.

### Storage Systems

Displays a list of all managed storage systems including their health status.

### Fabrics

Displays a list of all managed fabrics including their health status.

### Switches

Displays a list of all managed switches including their health status.

## Tape libraries

Displays a list of all managed tape libraries including their health status.

## TPC Server Overview workspace

The TPC Server Overview workspace provides an overview of several metrics about the Tivoli Storage Productivity Center server. You can access this workspace by clicking **TPC Server Overview** in the navigation tree of the Navigator window.

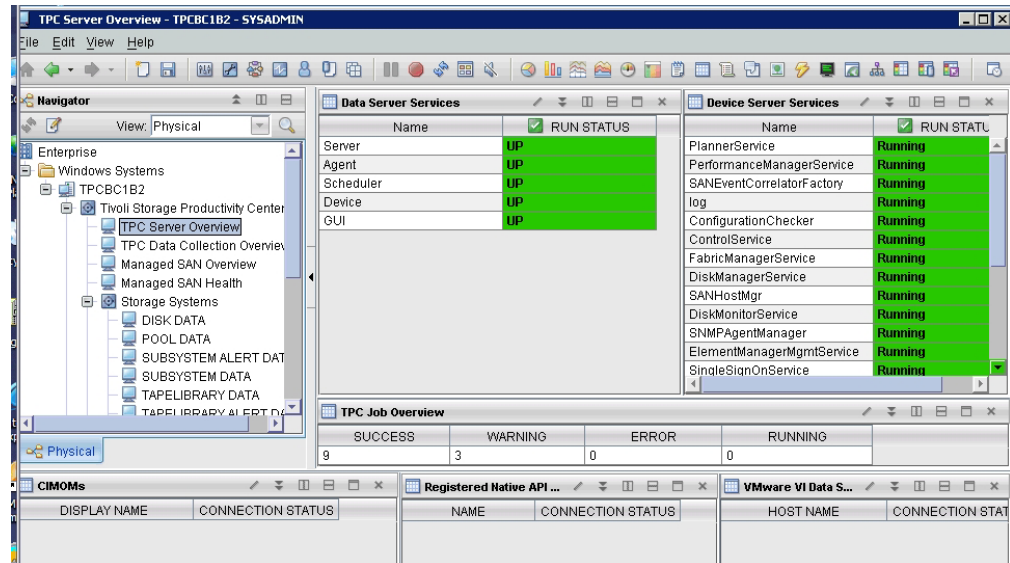


Figure 63. TPC Server overview report

This workspace provides the following views:

### Data Server Services

Displays a table showing all the Data server services (Server, Agent, Scheduler, Device, and GUI) including their run status.

### Device Server Services

Displays a table showing all Device server services including their run status.

### TPC Job Overview

Displays a table showing statistics about the total amount of jobs and their corresponding status.

### CIMOMs

Displays a table showing all the managed CIMOMs including their connection status.

### Registered Native API Devices

Displays a table showing all managed native interface devices including their corresponding connection status.

### VMware VI Data Sources

Displays a report showing all attached VMware VI Data Sources and their corresponding connection status.

## TPC Data Collection Overview workspace

The TPC Data Collection Overview workspace provides the status of the latest probe jobs and performance monitor jobs. You can access this workspace by clicking **TPC Data Collection Overview** in the navigation tree of the Navigator window.

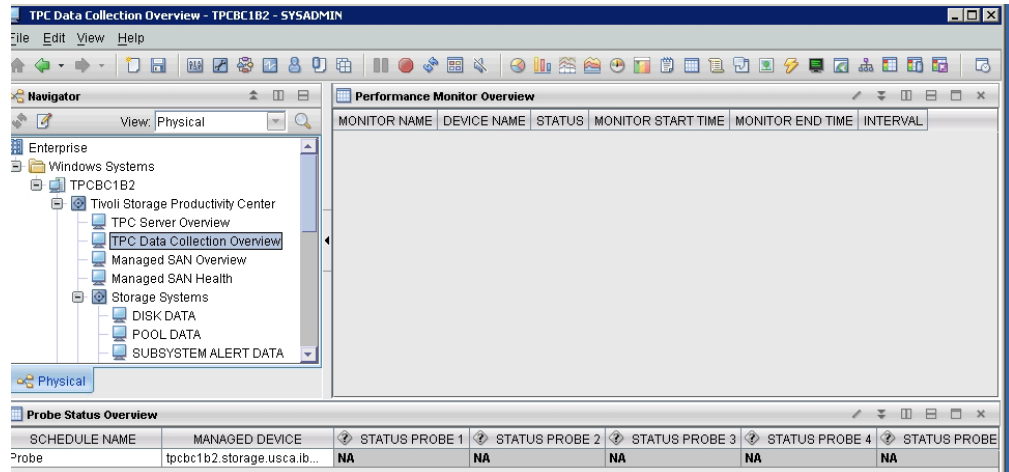


Figure 64. TPC Server overview report

This workspace provides the following information:

### Probe status overview

Displays a table showing all probe jobs and their status with the last five runs.

### Performance monitor overview

Displays a table showing all performance monitor jobs and their corresponding status.

## Detailed Tivoli Storage Productivity Center reports

There are several reports which provide detailed information about the monitored entities. You can access these reports by clicking the appropriate entity in the navigation tree of the Navigator window. These entities are:

### Storage Systems

Under **Storage Systems**, you can view these detailed reports:

- DISK DATA
- POOL DATA
- SUBSYSTEM ALERT DATA
- SUBSYSTEM DATA
- TAPELIBRARY DATA
- TAPELIBRARY ALERT DATA
- VMWARE VI DATA
- VOLUME DATA

### TPC Server and Data Sources

Under **TPC Server and Data Sources**, you can view these detailed reports:

- CIMOM DATA
- DATASERVER API

- DEVICESERVER API
- JOB STATUS DATA
- REGISTERED NATIVE API DATA
- OTHER DEVICE DATA
- PERFORMANCE MONITOR DATA
- PROBE STATUS DATA
- SCHEDULED JOB DATA
- TPC SERVER DATA

#### **Computers and Hypervisors**

Under **Computers and Hypervisors**, you can view these detailed reports:

- COMPUTER ALERT DATA
- COMPUTER DATA
- DATA AGENT DATA
- HBA DATA
- HYPERVISOR DATA
- INBANDAGENT DATA

#### **Fabrics and Switches**

Under **Fabrics and Switches**, you can view these detailed reports:

- FABRIC DATA
- FABRIC ALERT DATA
- OUTBANDAGENT DATA
- SWITCH DATA
- SWITCH ALERT DATA

See the Tivoli Storage Productivity Center online help for information about these entities.



---

## Appendix A. Reference

This topic contains reference material.

---

### Alerts and schedules associated with monitoring groups and templates

Use alerts to notify you when conditions for storage devices are detected by a data collection schedule. Include storage devices in monitoring groups or templates to automatically associate alerts and schedules with those devices.

IBM Tivoli Storage Productivity Center includes predefined groups and templates that are associated with different sets of alerts and schedules. The group or template that you select for a storage device determines what conditions are checked when data is collected for those devices. Use the information in this topic to help you determine the groups and templates that you want to use for the storage devices in your environment.

**Notes:**

- If you want to include a storage device in a group, you can select a user-defined group rather than a predefined group. Any alerts and schedules associated with that user-defined group are automatically applied to the devices.
- When not using the Configure Devices wizard, you are not required to include a storage device in a group or template to collect data or associate alerts with it. You can include storage devices individually in schedules and associate alerts with them at any time. Monitoring groups and templates are intended to help you assign the same schedules and alerts to a number of related devices at the same time.
- You can edit predefined monitoring groups and their associated probes, pings, and alerts at any time. You can also edit the probes, pings, and alerts that are created based on templates.

## Storage Subsystems

Table 52. Predefined monitoring groups, templates, schedules, and alerts for storage subsystems

| Group name / Template name                                                                          | The data collection schedule associated with a group or template                                                                                                                                                                                                 | Alert conditions that are checked for devices in the group or template during data collection <sup>1, 2</sup>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>Group:</b> Subsystem Basic Group</p> <p><b>Template:</b> Basic Subsystem Monitoring</p>       | <p><b>Schedule name:</b> Subsystem Basic Probe</p> <p><b>When it is run:</b> 1x/week on Tuesday at 1:00 am</p> <p>The schedule name and time are different for templates. See <sup>3</sup> at the bottom of this table for more information.</p>                 | <p>Subsystem Discovered<br/>           Disk Not Found<br/>           Storage Subsystem Not Found<br/>           Subsystem Status Change Offline<br/>           Subsystem Status Change Online<br/>           Volume Status Change Offline<br/>           Volume Status Change Online<br/>           Pool Status Change<br/>           Pool Discovered<br/>           Pool Status Change Offline<br/>           Pool Status Change Online<br/>           Volume Not Found<br/>           File System Is Missing<br/>           Cluster Was Rediscovered<br/>           Cluster Was Removed<br/>           NSD Is Missing<br/>           NSD State Changed</p>                                                                                                                                                                                                                                                                                                       |
| <p><b>Group:</b> Subsystem Standard Group</p> <p><b>Template:</b> Standard Subsystem Monitoring</p> | <p><b>Schedule name:</b> Subsystem Standard Probe</p> <p><b>When it is run:</b> 2x/week on Monday and Wednesday at 1:00 am</p> <p>The schedule name and time are different for templates. See <sup>3</sup> at the bottom of this table for more information.</p> | <p>Subsystem Discovered<br/>           Disk Not Found<br/>           Storage Subsystem Not Found<br/>           Subsystem Status Change Offline<br/>           Subsystem Status Change Online<br/>           Backend Controller Status Change Offline<br/>           Backend Controller Status Change Online<br/>           Volume Status Change Offline<br/>           Volume Status Change Online<br/>           Volume Capacity Change<br/>           Pool Status Change<br/>           Pool Discovered<br/>           Pool Status Change Offline<br/>           Pool Status Change Online<br/>           Pool Capacity Change<br/>           Volume Not Found<br/>           File System Is Missing<br/>           File System Reconfigured<br/>           Cluster Was Rediscovered<br/>           Cluster Was Removed<br/>           Export Is Missing<br/>           Export State Changed<br/>           NSD Is Missing<br/>           NSD State Changed</p> |



Table 52. Predefined monitoring groups, templates, schedules, and alerts for storage subsystems (continued)

| Group name / Template name                                                                          | The data collection schedule associated with a group or template                                                                                                                                                                         | Alert conditions that are checked for devices in the group or template during data collection <sup>1, 2</sup>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>Group:</b> Subsystem Advanced Group</p> <p><b>Template:</b> Advanced Subsystem Monitoring</p> | <p><b>Schedule name:</b> Subsystem Advanced Probe</p> <p><b>When it is run:</b> 1x/day at 10:00 pm</p> <p>The schedule name and time are different for templates. See <sup>3</sup> at the bottom of this table for more information.</p> | <p>Subsystem Discovered<br/>           Cache Increased<br/>           Cache Decreased<br/>           Disk Not Found<br/>           Storage Subsystem Not Found<br/>           Subsystem Status Change Offline<br/>           Subsystem Status Change Online<br/>           Subsystem Version Change<br/>           Subsystem Property Change<br/>           Subsystem Allocated Capacity Change<br/>           Subsystem Available Capacity Change<br/>           Subsystem Backend Capacity Change<br/>           Backend Controller Status Change Offline<br/>           Backend Controller Status Change Online<br/>           Volume Status Change Offline<br/>           Volume Status Change Online<br/>           Volume Capacity Change<br/>           Pool Status Change<br/>           Pool Discovered<br/>           Pool Status Change Offline<br/>           Pool Status Change Online<br/>           Pool Capacity Change<br/>           Node Status Change Offline<br/>           Node Status Change Online<br/>           Node Status Change<br/>           Node Discovered<br/>           Datapath State Change<br/>           Volume Not Found<br/>           File System Is Missing<br/>           File System Reconfigured<br/>           Cluster Was Rediscovered<br/>           Cluster Was Removed<br/>           Export Is Missing<br/>           Export Path Changed<br/>           Export Protocol Changed<br/>           Export State Changed<br/>           NSD Is Missing<br/>           NSD State Changed<br/>           Node State Changed<br/>           Node Status Changed to Error<br/>           Node Status Changed to Offline<br/>           Node Status Changed to Online<br/>           Node Status Changed to Warning</p> |
| <p>Default Storage Subsystem Group<sup>4</sup></p>                                                  | <p>No schedule is associated with this group</p>                                                                                                                                                                                         | <p>No alert conditions are associated with this group</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

Table 52. Predefined monitoring groups, templates, schedules, and alerts for storage subsystems (continued)

| Group name / Template name                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | The data collection schedule associated with a group or template | Alert conditions that are checked for devices in the group or template during data collection <sup>1, 2</sup> |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------|
| <p><b>Notes:</b></p> <ol style="list-style-type: none"> <li>See Triggering conditions for storage subsystem alerts for descriptions of the alert triggering conditions for storage subsystems.</li> <li>Performance alerts for storage subsystems are not associated with groups and templates. Performance alerts are alerts that contain a "threshold" triggering condition. If you want to include a storage subsystem in a performance alert, you must create an alert with a threshold triggering condition and then individually include the storage subsystem in that alert. See "Creating storage system alerts" on page 390 for information about how to create storage subsystem alerts.</li> <li>Note the following differences between schedules created from templates and groups: <ul style="list-style-type: none"> <li>The name of a probe schedule created from a template contains a prefix that is defined in the Configure Devices wizard. For example, if the Basic Subsystem Monitoring template is selected and "ABC" is entered as the prefix, the resulting schedule is named "<i>userID</i>.ABCssBasicProbe", where <i>userID</i> represents the ID of the user that is logged in to Tivoli Storage Productivity Center.</li> <li>The probe schedule created from a template might not run at the time specified in this table. Tivoli Storage Productivity Center determines the time at which the probe is run to ensure optimal server performance.</li> </ul> </li> <li>The Default Storage Subsystem Group is not automatically associated with a schedule or alerting conditions. However, you can manually include this group in a schedule or associate it with alerts at any time.</li> </ol> |                                                                  |                                                                                                               |

## Fabrics

Table 53. Predefined monitoring groups, templates, schedules, and alerts for fabrics

| Group name / Template name                                                                    | The data collection schedule associated with a group or template                                                                                                                                                                                           | Alert conditions that are checked for devices in the group or template during data collection <sup>1</sup>                                                                                                                                                                                                            |
|-----------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>Group:</b> Fabric Basic Group</p> <p><b>Template:</b> Basic Fabric Monitoring</p>       | <p><b>Schedule name:</b> Fabric Basic Probe</p> <p><b>When it is run:</b> 2x/week on Monday and Wednesday at 3:00 am</p> <p>The schedule name and time are different for templates. See <sup>2</sup> at the bottom of this table for more information.</p> | <p>Fabric State Change<br/> Fabric to Switch Change<br/> Switch discovered<br/> Switch State Change<br/> Switch Status Change Offline<br/> Switch Status Change Online</p>                                                                                                                                            |
| <p><b>Group:</b> Fabric Standard Group</p> <p><b>Template:</b> Standard Fabric Monitoring</p> | <p><b>Schedule name:</b> Fabric Standard Probe</p> <p><b>When it is run:</b> 1x/day at 4:00 am</p> <p>The schedule name and time are different for templates. See <sup>2</sup> at the bottom of this table for more information.</p>                       | <p>Fabric State Change<br/> Fabric to Switch Change<br/> New Connection<br/> Connection State Change<br/> Switch discovered<br/> Switch State Change<br/> Switch Status Change Offline<br/> Switch Status Change Online<br/> Switch Blade Change<br/> Switch Blade Change Offline<br/> Switch Blade Change Online</p> |

Table 53. Predefined monitoring groups, templates, schedules, and alerts for fabrics (continued)

| Group name / Template name                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | The data collection schedule associated with a group or template                                                                                                                                                                          | Alert conditions that are checked for devices in the group or template during data collection <sup>1</sup>                                                                                                                                                                                                                                                                                                                                         |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Group:</b> Fabric Advanced Group<br><br><b>Template:</b> Advanced Fabric Monitoring                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>Schedule name:</b> Fabric Advanced Probe<br><br><b>When it is run:</b> 2x/day at 6:00 am and 6:00 pm<br><br>The schedule name and time are different for templates. See <sup>2</sup> at the bottom of this table for more information. | Fabric State Change<br>Fabric to Switch Change<br>Zone State Change<br>Zone Set State Change<br>Zone Set Discovered<br>Zone Set to Zone Change<br>Zone Alias to Member Change<br>New Connection<br>Connection State Change<br>Switch discovered<br>Switch State Change<br>Switch Status Change Offline<br>Switch Status Change Online<br>Switch Version Change<br>Switch Blade Change<br>Switch Blade Change Offline<br>Switch Blade Change Online |
| Default Fabric Group <sup>3</sup>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | No schedule is associated with this group                                                                                                                                                                                                 | No alert conditions are associated with this group                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Notes:</b><br>1. See Triggering conditions for fabric alerts for descriptions of the alert triggering conditions for fabrics. See Triggering conditions for switch alerts for descriptions of the alert triggering conditions for switches.<br>2. Note the following differences between schedules created from templates and groups: <ul style="list-style-type: none"> <li>• The name of a probe schedule created from a template contains a prefix that is defined in the Configure Devices wizard. For example, if the Basic Fabric Monitoring template is selected and "ABC" is entered as a prefix, the resulting schedule is named "<i>userID</i>.ABCfabBasicProbe", where <i>userID</i> represents the ID of the user that is logged in to Tivoli Storage Productivity Center.</li> <li>• The probe schedule created from a template might not run at the time specified in this table. Tivoli Storage Productivity Center determines the time at which the probe is run to ensure optimal server performance.</li> </ul> 3. The Default Fabric Group is not automatically associated with a schedule or alerting conditions. However, you can manually include this group in a schedule or associate it with alerts at any time. |                                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## Computers

Table 54. Predefined monitoring groups, templates, schedules, and alerts for computers

| Group name / Template name                                                           | The data collection schedule associated with a group or template                                                                                                                                                                                                                                                                                                                                                     | Alert conditions that are checked for devices in the group or template during data collection <sup>1</sup>                                                                          |
|--------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Group:</b> Computer Basic Group<br><br><b>Template:</b> Basic Computer Monitoring | <ul style="list-style-type: none"> <li>• Ping schedule<br/> <b>Name:</b> Computer Basic Ping<br/> <b>When it is run:</b> every 3 hours</li> <li>• Probe schedule<br/> <b>Name:</b> Computer Basic Probe<br/> <b>When it is run:</b> 2x/week on Monday and Wednesday at 3:00 am</li> </ul> The schedule name and time are different for templates. See <sup>2</sup> at the bottom of this table for more information. | Disk Not Found<br>New Disk Defect Found<br>Disk Failure Predicted<br>Computer Unreachable<br>Computer Discovered<br>Computer Status Change Offline<br>Computer Status Change Online |

Table 54. Predefined monitoring groups, templates, schedules, and alerts for computers (continued)

| Group name / Template name                                                                        | The data collection schedule associated with a group or template                                                                                                                                                                                                                                                                                                                                                   | Alert conditions that are checked for devices in the group or template during data collection <sup>1</sup>                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>Group:</b> Computer Standard Group</p> <p><b>Template:</b> Standard Computer Monitoring</p> | <ul style="list-style-type: none"> <li>• Ping schedule<br/><b>Name:</b> Computer Standard Ping<br/><b>When it is run:</b> every 60 minutes</li> <li>• Probe schedule<br/><b>Name:</b> Computer Standard Probe<br/><b>When it is run:</b> 1x/day at 3:00 am</li> </ul> <p>The schedule name and time are different for templates. See <sup>2</sup> at the bottom of this table for more information.</p>            | <p>New Disk Detected<br/>Disk Not Found<br/>New Disk Defect Found<br/>Grown Disk Defects Exceed 5<br/>Disk Failure Predicted<br/>New Filesystem Detected<br/>Computer Unreachable<br/>Computer Discovered<br/>Computer Status Change Offline<br/>Computer Status Change Online<br/>HBA Driver Version Change<br/>HBA Firmware Version Change</p>                                                                                                                                                                                                                                 |
| <p><b>Group:</b> Computer Advanced Group</p> <p><b>Template:</b> Advanced Computer Monitoring</p> | <ul style="list-style-type: none"> <li>• Ping schedule<br/><b>Name:</b> Computer Advanced Ping<br/><b>When it is run:</b> every 30 minutes</li> <li>• Probe schedule<br/><b>Name:</b> Computer Advanced Probe<br/><b>When it is run:</b> 2x/day at 5:00 am at 5:00 pm</li> </ul> <p>The schedule name and time are different for templates. See <sup>2</sup> at the bottom of this table for more information.</p> | <p>RAM Increased<br/>RAM Decreased<br/>Virtual Memory Increased<br/>Virtual Memory Decreased<br/>New Disk Detected<br/>Disk Not Found<br/>New Disk Defect Found<br/>Grown Disk Defects Exceed 5<br/>Disk Failure Predicted<br/>New Filesystem Detected<br/>Cluster Resource Group Added<br/>Cluster Resource Group Removed<br/>Cluster Resource Group Moved<br/>Computer Unreachable<br/>Computer Discovered<br/>Computer Status Change Offline<br/>Computer Status Change Online<br/>Computer Property Change<br/>HBA Driver Version Change<br/>HBA Firmware Version Change</p> |
| <p>Default Computer Group<sup>3</sup></p>                                                         | <ul style="list-style-type: none"> <li>• Ping schedule<br/><b>Name:</b> Default Ping<br/><b>When it is run:</b> every 30 minutes</li> <li>• Probe schedule<br/><b>Name:</b> Default Probe<br/><b>When it is run:</b> 1x/day at 1:45 am</li> </ul>                                                                                                                                                                  | <p>No alert conditions are associated with this group</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

**Notes:**

1. See Triggering conditions for computer alerts for descriptions of the alert triggering conditions for computers.
2. Note the following differences between schedules created from templates and groups:
  - The name of a probe or ping schedule created from a template contains a prefix that is defined in the Configure Devices wizard. For example, if the Basic Computer Monitoring template is selected and "ABC" is entered as a prefix, the resulting schedule is named "*userID*.ABCcompBasicProbe", where *userID* represents the ID of the user that is logged in to Tivoli Storage Productivity Center.
  - The probe schedule created from a template might not run at the time specified in this table. Tivoli Storage Productivity Center determines the time at which the probe is run to ensure optimal server performance.
3. The Default Computer Group is not automatically associated with alerting conditions. However, you can manually associate it with alerts at any time.

## Tape Libraries

Table 55. Predefined monitoring groups, templates, and schedules for tape libraries<sup>1</sup>

| Group name / Template name                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | The data collection schedule associated with a group or template                                                                                                                                                                             |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Group:</b> Tape Basic Group<br><b>Template:</b> Basic Tape Monitoring                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>Schedule name:</b> Tape Basic Probe<br><b>When it is run:</b> 1x/week on Sunday at 4:00 am<br>The schedule name and time are different for templates. See <sup>2</sup> at the bottom of this table for more information.                  |
| <b>Group:</b> Tape Standard Group<br><b>Template:</b> Standard Tape Monitoring                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>Schedule name:</b> Tape Standard Probe<br><b>When it is run:</b> 2x/week on Sunday and Wednesday at 4:00 am<br>The schedule name and time are different for templates. See <sup>2</sup> at the bottom of this table for more information. |
| <b>Group:</b> Tape Advanced Group<br><b>Template:</b> Advanced Tape Monitoring                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>Schedule name:</b> Tape Advanced Probe<br><b>When it is run:</b> 1x/day at 6:00 am<br>The schedule name and time are different for templates. See <sup>2</sup> at the bottom of this table for more information.                          |
| Default Tape LibraryGroup <sup>2</sup>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | No schedule is associated with this group                                                                                                                                                                                                    |
| <b>Notes:</b> <ol style="list-style-type: none"> <li>There are no alerting conditions available for tape libraries.</li> <li>Note the following differences between schedules created from templates and groups:             <ul style="list-style-type: none"> <li>The name of a probe schedule created from a template contains a prefix that is defined in the Configure Devices wizard. For example, if the Basic Tape Monitoring template is selected and "ABC" is entered as a prefix, the resulting schedule is named "<i>userID</i>.ABCTapeBasicProbe", where <i>userID</i> represents the ID of the user that is logged in to Tivoli Storage Productivity Center.</li> <li>The probe schedule created from a template might not run at the time specified in this table. Tivoli Storage Productivity Center determines the time at which the probe is run to ensure optimal server performance.</li> </ul> </li> <li>The Default Tape Library Group is not automatically associated with a schedule. However, you can manually include this group in a schedule at any time.</li> </ol> |                                                                                                                                                                                                                                              |

## Supported Fabric Agent Types

The following table shows the supported agent types for Switch and Fabric functions.

Table 56. Supported agent types for Switch and Fabric Functions

| Function > Switch             | Brocade     | McDATA                                                                                                                               | Cisco <sup>1</sup>                                                                         | QLogic                                                            |
|-------------------------------|-------------|--------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|-------------------------------------------------------------------|
| Switch Performance Monitoring | CIMOM agent | CIMOM agent                                                                                                                          | CIMOM agent                                                                                | CIMOM agent                                                       |
| Zone Control <sup>2</sup>     | CIMOM agent | Recommended:<br>CIMOM agent <sup>6</sup><br><br>Also supported:<br>Storage Resource agent<br>or In-band Fabric agent<br><sup>6</sup> | Storage Resource agent<br>or In-band Fabric agent<br>required in each VSAN<br><sup>6</sup> | Storage Resource agent<br>or In-band Fabric agent<br><sup>6</sup> |

Table 56. Supported agent types for Switch and Fabric Functions (continued)

| Function > Switch                                                              | Brocade                                                                                                                        | McDATA                                                                                                                         | Cisco <sup>1</sup>                                                                                                             | QLogic                                                                   |
|--------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------|
| Zone Control with Zone Aliases                                                 | CIMOM agent                                                                                                                    | Not supported                                                                                                                  | Not supported                                                                                                                  | Not supported                                                            |
| Switch and switch port information collected                                   | Recommended:<br>CIMOM agent<br><br>Also supported:<br>Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent | Recommended:<br>CIMOM agent<br><br>Also supported:<br>Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent | Recommended:<br>CIMOM agent<br><br>Also supported:<br>Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent | Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent |
| Topology connectivity information collected                                    | Recommended:<br>CIMOM agent<br><br>Also supported:<br>Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent | Recommended:<br>CIMOM agent<br><br>Also supported:<br>Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent | Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent                                                       | Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent |
| Zoning information collected <sup>2</sup>                                      | Recommended:<br>CIMOM agent                                                                                                    | Recommended:<br>CIMOM agent<br><br>Also supported:<br>In-band Fabric agent or Storage Resource agent                           | In-band Fabric agent or Storage Resource agent required in each VSAN <sup>3 4 5</sup>                                          | In-band Fabric agent or Storage Resource agent                           |
| Tivoli Storage Productivity Center Alerts                                      | Recommended:<br>CIMOM agent<br><br>Also supported:<br>Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent | Recommended:<br>CIMOM agent<br><br>Also supported:<br>Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent | Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent                                                       | Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent |
| Hosts, endpoint devices, device-centric and host-centric information collected | In-band Fabric agent or Storage Resource agent                                                                                 | In-band Fabric agent or Storage Resource agent                                                                                 | In-band Fabric agent or Storage Resource agent                                                                                 | In-band Fabric agent or Storage Resource agent                           |
| Switch "Sensors and Events"                                                    | Out-of-band Fabric agent                                                                                                       | Out-of-band Fabric agent                                                                                                       | Out-of-band Fabric agent                                                                                                       | Out-of-band Fabric agent                                                 |

Table 56. Supported agent types for Switch and Fabric Functions (continued)

| Function > Switch                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Brocade | McDATA | Cisco <sup>1</sup> | QLogic |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------|--------------------|--------|
| <p><b>Notes:</b></p> <p><sup>1</sup> The In-band Fabric agent or Storage Resource agent gets the information on a VSAN basis. Each VSAN is viewed as an individual SAN. The Out-of-band Fabric agent and CIMOM agent get the physical fabric information and can correlate the VSAN information to a physical infrastructure. The Out-of-band Fabric agent also collects some VSAN information.</p> <p><sup>2</sup> Zone Control refers to the ability to actively configure zoning. "Zone information collected" indicates that the active zone set and the inactive zoning library information is collected during a fabric probe.</p> <p><sup>3</sup> For many Cisco proprietary zone members, the values are displayed as zeroes.</p> <p><sup>4</sup> If a zoning configuration has Cisco proprietary zone members, then you will not be able to modify the zoning for that configuration through Tivoli Storage Productivity Center. You need to use the element manager to remove proprietary zone members or do zoning.</p> <p><sup>5</sup> Orphan Cisco zones are not reported.</p> <p><sup>6</sup> Tivoli Storage Productivity Center will not get a transaction lock on the switch while making zone changes on the fabric.</p> |         |        |                    |        |

## Supported storage systems providing full disk encryption and solid-state drives

IBM Tivoli Storage Productivity Center supports full disk encryption and solid-state drives in the IBM System Storage DS8000 and IBM System Storage DS5000 Storage Manager systems described in this topic.

### DS8000 4.2 and later

Tivoli Storage Productivity Center supports full disk encryption and solid-state drives in DS8000 4.2 and later.

### IBM System Storage DS5000 Storage Manager series

IBM Tivoli Storage Productivity Center supports full disk encryption and solid-state drives in the following DS5000 series systems:

- DS5100 and DS5300 - full disk encryption, solid-state drives, 1 TB SATA drives
- DS5020 - full disk encryption

To support full disk encryption and solid-state drives in DS5000, the LSI SMI-S Provider is required. For more information about the LSI SMI-S Provider, see the LSI Web site at <http://www.lsi.com>.

## Frequently Asked Questions

Use this section to view answers to common questions about IBM Tivoli Storage Productivity Center. These questions are organized according to the manager to which they apply.

### General and Migration Information

Learn about the issues when working with IBM Tivoli Storage Productivity Center, upgrading to a new version of the product, or migrating from another product.

#### What Java version is officially supported by Tivoli Storage Productivity Center version 4.1 graphical user interface?

Java 5 is the level supported for the Tivoli Storage Productivity Center

GUI. For information about how to start the Tivoli Storage Productivity Center graphical user interface, see the information center. Search for **Starting and stopping the console**.

**You are entering namespaces manually for CIM agents and switches in the Tivoli Storage Productivity Center user interface. What are the current namespaces for the supported brands?**

Table 57. Interop Namespaces for CIM Agents for switches and storage systems

| Switch or system                                        | Namespace                                                                                                                                                                                                                                                                                                                         |
|---------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IBM (other than IBM System Storage TS3310 Tape Library) | /root/ibm                                                                                                                                                                                                                                                                                                                         |
| IBM System Storage TS3310 Tape Library                  | /root/cimv2                                                                                                                                                                                                                                                                                                                       |
| Brocade                                                 | /interop                                                                                                                                                                                                                                                                                                                          |
| Cisco                                                   | /root/cimv2<br><br>For version 3.2.1 or later:<br>/root/pg_interop                                                                                                                                                                                                                                                                |
| Engenio                                                 | /interop                                                                                                                                                                                                                                                                                                                          |
| EMC                                                     | /root/emc                                                                                                                                                                                                                                                                                                                         |
| HDS                                                     | For HiCommand 5.6 or later, use: /root/smis/current<br><br>For a HiCommand version that supports SMI-S 1.2, use following namespace to traverse the model as SMI-S 1.1:<br>/root/smis/smis11<br><br>For HiCommand versions earlier than HiCommand 5.6, use:<br>/root/hitachi/dmxx, where xx is represents the level of HiCommand. |
| HP                                                      | /root                                                                                                                                                                                                                                                                                                                             |
| LSI                                                     | /root/PG_InterOp                                                                                                                                                                                                                                                                                                                  |
| McData                                                  | /interop                                                                                                                                                                                                                                                                                                                          |
| SUN                                                     | /root/sun3510 or /interop<br><br>Note: This namespace is for a system and not a switch.                                                                                                                                                                                                                                           |
| XVRATEX                                                 | /root/PG_interop                                                                                                                                                                                                                                                                                                                  |

Check with your switch vendor to ensure you use the most current namespaces.

**What if you specify an incorrect namespace?**

If you specify an incorrect namespace the following issues might occur:

- The connection test fails when the CIMOM is added.
- The discovery does not discover all information of the system managed by the CIMOM.
- The probe fails.
- The function you want to perform on the system might fail (for example, collecting performance data).

**What are the required steps when upgrading the Tivoli Storage Productivity Center to the latest version?**

See the Tivoli Storage Productivity Center information center for detailed information about how to upgrade the product:<http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp>. To access



information about upgrading, expand the **Tivoli Storage Productivity Center > Upgrading** node in the information center.

**You are using an older version of IBM Tivoli Storage Productivity Center for Data. How can I upgrade to this latest version?**

See the Tivoli Storage Productivity Center information center for detailed information about how to upgrade the product: <http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp>. To access information about upgrading, expand the **Tivoli Storage Productivity Center > Upgrading** node in the navigation tree.

## Data Manager

Use the questions in this section to answer questions about working with the Data Manager component of IBM Tivoli Storage Productivity Center.

**How do you know if your storage system is supported by Data Manager and which SMI-S agents are supported?**

To confirm whether your storage system is supported and which SMI-S agent is supported for that system, review the supported products list for the current release of Tivoli Storage Productivity Center. Go to [http://www-01.ibm.com/support/docview.wss?rs=40&context=SWJ50&q1=matrix&uid=swg21386446&loc=en\\_US&cs=utf-8&lang=en](http://www-01.ibm.com/support/docview.wss?rs=40&context=SWJ50&q1=matrix&uid=swg21386446&loc=en_US&cs=utf-8&lang=en) and click the applicable Tivoli Storage Productivity Center release in the Storage column.

## Fabric Manager

Use the questions in this section to answer questions about working with the Fabric Manager component of IBM Tivoli Storage Productivity Center.

Keep in mind the following information and considerations:

**What type of fabric agents should you be using in Tivoli Storage Productivity Center?**

The following table shows the current agent types supported for the Switch Performance Management functionality and the Fabric Zone Configuration functionality.

*Table 58. Supported agent types for Switch and Fabric Functions*

| Function > Switch              | Brocade     | McDATA                                                                                                                               | Cisco <sup>1</sup>                                                                         | QLogic                                                            |
|--------------------------------|-------------|--------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|-------------------------------------------------------------------|
| Switch Performance Monitoring  | CIMOM agent | CIMOM agent                                                                                                                          | CIMOM agent                                                                                | CIMOM agent                                                       |
| Zone Control <sup>2</sup>      | CIMOM agent | Recommended:<br>CIMOM agent <sup>6</sup><br><br>Also supported:<br>Storage Resource agent<br>or In-band Fabric agent<br><sup>6</sup> | Storage Resource agent<br>or In-band Fabric agent<br>required in each VSAN<br><sup>6</sup> | Storage Resource agent<br>or In-band Fabric agent<br><sup>6</sup> |
| Zone Control with Zone Aliases | CIMOM agent | Not supported                                                                                                                        | Not supported                                                                              | Not supported                                                     |

Table 58. Supported agent types for Switch and Fabric Functions (continued)

| Function > Switch                                                              | Brocade                                                                                                                        | McDATA                                                                                                                         | Cisco <sup>1</sup>                                                                                                             | QLogic                                                                   |
|--------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------|
| Switch and switch port information collected                                   | Recommended:<br>CIMOM agent<br><br>Also supported:<br>Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent | Recommended:<br>CIMOM agent<br><br>Also supported:<br>Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent | Recommended:<br>CIMOM agent<br><br>Also supported:<br>Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent | Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent |
| Topology connectivity information collected                                    | Recommended:<br>CIMOM agent<br><br>Also supported:<br>Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent | Recommended:<br>CIMOM agent<br><br>Also supported:<br>Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent | Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent                                                       | Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent |
| Zoning information collected <sup>2</sup>                                      | Recommended:<br>CIMOM agent                                                                                                    | Recommended:<br>CIMOM agent<br><br>Also supported:<br>In-band Fabric agent or Storage Resource agent                           | In-band Fabric agent or Storage Resource agent required in each VSAN <sub>3 4 5</sub>                                          | In-band Fabric agent or Storage Resource agent                           |
| Tivoli Storage Productivity Center Alerts                                      | Recommended:<br>CIMOM agent<br><br>Also supported:<br>Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent | Recommended:<br>CIMOM agent<br><br>Also supported:<br>Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent | Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent                                                       | Out-of-band Fabric agent, In-band Fabric agent or Storage Resource agent |
| Hosts, endpoint devices, device-centric and host-centric information collected | In-band Fabric agent or Storage Resource agent                                                                                 | In-band Fabric agent or Storage Resource agent                                                                                 | In-band Fabric agent or Storage Resource agent                                                                                 | In-band Fabric agent or Storage Resource agent                           |
| Switch "Sensors and Events"                                                    | Out-of-band Fabric agent                                                                                                       | Out-of-band Fabric agent                                                                                                       | Out-of-band Fabric agent                                                                                                       | Out-of-band Fabric agent                                                 |

Table 58. Supported agent types for Switch and Fabric Functions (continued)

| Function > Switch                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Brocade | McDATA | Cisco <sup>1</sup> | QLogic |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------|--------------------|--------|
| <b>Notes:</b><br><sup>1</sup> The In-band Fabric agent or Storage Resource agent gets the information on a VSAN basis. Each VSAN is viewed as an individual SAN. The Out-of-band Fabric agent and CIMOM agent get the physical fabric information and can correlate the VSAN information to a physical infrastructure. The Out-of-band Fabric agent also collects some VSAN information.<br><sup>2</sup> Zone Control refers to the ability to actively configure zoning. "Zone information collected" indicates that the active zone set and the inactive zoning library information is collected during a fabric probe.<br><sup>3</sup> For many Cisco proprietary zone members, the values are displayed as zeroes.<br><sup>4</sup> If a zoning configuration has Cisco proprietary zone members, then you will not be able to modify the zoning for that configuration through Tivoli Storage Productivity Center. You need to use the element manager to remove proprietary zone members or do zoning.<br><sup>5</sup> Orphan Cisco zones are not reported.<br><sup>6</sup> Tivoli Storage Productivity Center will not get a transaction lock on the switch while making zone changes on the fabric. |         |        |                    |        |

### What are the prerequisites for using the Switch Performance functionalities in Tivoli Storage Productivity Center?

The prerequisite steps for using the Switch Performance functionalities are:

1. Add a CIMOM as a data source to the Tivoli Storage Productivity Center system that manages the switch.
2. Run a CIMOM discovery and wait for it to complete successfully for that CIMOM.
3. Run a probe against the fabric and wait for it to complete successfully.

After you complete these steps, create a Switch Performance monitor to trigger defined switch performance alert conditions or generate switch performance reports.

### Why does an Out of Band Fabric (SNMP) agent show a state of "not contacted"?

Perform the following tasks to help determine why an Out of Band Fabric (SNMP) agent show a state of "not contacted":

- Check that the IP Address or name is entered correctly in the Out of Band Fabric agent details panel.
- Ensure the FA MIB is enabled on the switch.
- Ensure that the Tivoli Storage Productivity Center server IP address is included in any SNMP Access Control List that exists on the switches.
- Ensure you are using SNMP v1.
- Ensure the string on the Tivoli Storage Productivity Center Out of Band Fabric agent details panel in the SNMP Community field matches one of the community strings on the switch.

### Can you add a storage subsystem as an Out of Band Fabric (SNMP) agent?

No. Out of Band Fabric agents within Tivoli Storage Productivity Center are meant for switches only. Disk Manager uses CIMOMs to gather subsystem information.

### Why aren't all the properties of a switch populated in the Topology Switch views?

A combination of Inband and CIMOM agents or a combination of Inband and Out of Band fabric agents might be required to gather all topology information for a switch. However, all of the properties are not normally needed.

**How can you specify a different management application to launch for a switch?**

Launch the Element Manager for a switch from the topology viewer. To do this, access the Switch view, right click a switch, and select **Launch Element Manager**. To change what is launched, right click a switch and select the **Launch Details Panel**. The element manager information is displayed at the bottom of the Details tab. To change this information, press **Set** and enter the HTTP, Telnet, or application information into the pop-up dialog and press **OK**.

**How do you manually change a fabric or switch label that is displayed throughout the system?**

Access the topology viewer in the Fabric or Switch view as appropriate. Right-click the entity you want to label to and select **Launch Details Panel**. Within the Details Panel, change the data in the **Label** field and click the save icon (or press Ctrl S).

**Why do you get the error "Agent capable of configuration zoning could not be found on this Fabric" for a fabric?**

For Brocade switches you must have a CIMOM agent configured for this fabric. A CIMOM agent is used by default if present. If a CIMOM for the fabric is present, you must run a CIMOM discovery and probe of the fabric before you can perform zone control.

For McDATA switches, a CIMOM agent is used by default (if available). If a CIMOM agent is not available, an Inband Fabric agent or Storage Resource Agent is used. If a CIMOM agent for the fabric is present, you must run a CIMOM discovery and probe of the fabric before you can perform zone control.

For all other supported switches, you must have a Tivoli Storage Productivity Center Inband Fabric agent in the fabric.

**Why is it not possible to configure inactive zoneset information for a McDATA fabric?**

If no CIMOM agent is configured for the McDATA fabric, and only Inband Fabric Agents or Storage Resource Agents are configured, then the zone control feature is limited to configuring only the active zoneset. Configure a CIMOM agent for the McDATA fabric to work with inactive zonesets for this fabric.

## **Disk Manager**

Use the questions in this section to answer questions about working with the Disk Manager component of IBM Tivoli Storage Productivity Center.

Keep in mind the following information and considerations:

**You encounter errors while collecting performance data on SAN Volume Controller. You fail to associate SAN Volume Controller performance data from non-configuration node with SAN Volume Controller performance data from configuration node. You encounter incomplete SAN Volume Controller performance data sample.**

This issue is caused by a configuration issue with SAN Volume Controller (time zone). Reset the time zone on SAN Volume Controller by logging into the SAN Volume Controller through putty. Run this command first:

```
svctask settimezone -timezone 509
```

This forces the cluster into the Universal time zone. To get the time zone you want the cluster to be in, run this command:

```
svctask settimezone -timezone
```

**Click on Launch Element Manager. The URL that appears is <http://0.0.0.0/ica>**

This error is caused by the wrong configuration of the SAN Volume Controller ICAT. Use the SAN Volume Controller GUI to remove and add the cluster again to the ICAT. This populates the correct value for the access point which is stored in the database repository during the next discovery. Use this cli command to set the IP address and port of the user interface.

For example: `svcluster -x modify -p 9.155.62.91:9080`

This is the same command used by the user interface and CIMOM.

---

## Arguments and window locations available with the inbound launch in context feature

This section describes the arguments and window locations available with the inbound launch in context feature.

### Launch in context parameters

You can use the launch in context parameters, definitions, and examples.

The command format for Java Web Start is unique in that a parameter and its value are separated by an equal sign (=). Parameter=value pairs are separated by an ampersand (&).

Not all parameters can be used with all launch types. Where this is true, it is noted in the parameter definition.

#### Parameters

**-user** The user ID. The value must be a valid user ID to log on to the IBM Tivoli Storage Productivity Center server, for example, `-user admin`.

**Note:** This parameter is not supported in a Web Browser's URL when using Java Web Start to launch the Tivoli Storage Productivity Center user interface.

**-passwd**

The password. The value must be a valid password to log on to the Tivoli Storage Productivity Center server, for example, `-passwd MyPassword`.

**Note:** This parameter is not supported in a Web Browser's URL when using Java Web Start to launch the Tivoli Storage Productivity Center user interface.

**-encpasswd**

The password encrypted. This is used if you choose to encrypt the password instead of entering the text password, for example, `-encpasswd 48735`.

**Tip:** To encrypt a password using Tivoli Storage Productivity Center (the preferred method), enter the following command in the Command Prompt window:

```
TPC-HOME\data/server/tools/tpctool encrypt text password
```

The encrypted password is returned.

**Note:** This parameter is not supported in a Web Browser's URL when using Java Web Start to launch the Tivoli Storage Productivity Center user interface.

**-encmethod**

The encryption method used to encrypt the password. The password can be encrypted using Tivoli Storage Productivity Center or Tivoli Enterprise Portal. The preferred method is to use Tivoli Storage Productivity Center to encrypt the password. If **encpasswd** is selected, you must also select **encmethod** with a value. The default is TPC. For example, `-encpasswd 48735 -encmethod TPC`

**Note:** This parameter is not supported in a Web Browser's URL when using Java Web Start to launch the Tivoli Storage Productivity Center user interface.

**Tip:** The Tivoli Enterprise Portal graphical user interface (GUI) does not start the Tivoli Storage Productivity Center GUI properly with a Tivoli Enterprise Portal-encrypted password if that encryption creates a string with spaces in it. Such a string is handled as multiple values unless it is enclosed in double quotation marks ("").

**Tip:** Encryption using Tivoli Enterprise Portal creates a string pattern. There are multiple string patterns that can result in the same decrypted value, so it is possible that different string patterns used as an encrypted password might result in same decrypted password.

**-function**

Specify one of the following values:

- The task to perform on the remote data server. See "Launch in context task parameters" on page 1153 for tasks and their specific parameters. For example, `-function disk.create_volume -subsystem_id <ID of the subsystem>`.
- The window name to open on the target system. The window name is mapped to the name of the node on the navigation tree. See "Available windows" on page 1154 for a list of available windows. For example, `-function topology.computer`.

**-server**

The Tivoli Storage Productivity Center server to log on to. This argument specifies a single server. For example, `-server tivoli11:9549`.

**-servers**

**Not used with the Web Browser.** The Tivoli Storage Productivity Center server or servers to log on to, in the format `server:port`. You can specify multiple server ports separated by commas. In that case, Tivoli Storage Productivity Center will connect to the first server. When a new connection from the top menu is created, (**Connection** > **New Connection**) the server list box includes the available servers in the **Connection** dialog. For example, `-servers "tivoli11:9549,wanda:9549,invader:9549"`

**-help**

**Not used with the Web Browser.** The help argument is used alone in the Arguments text field. It is entered as `-help functions`. This argument returns a list of the windows representing nodes in the navigation tree that can be launched. Typing `-help` with no arguments will return an error message.

You can also enter **help** as a command in the Command Prompt window as `tpc.bat -help functions` for Windows, and `TPCD.sh -help functions` for AIX, Linux, or UNIX.

## Launch in context task parameters

Use the launch in context feature task parameters to perform tasks.

Task parameters associated with actions taken on the Data server also use the `-function` parameter. For example:`-function data.filesystem_capacity_by_computer -computer_id "<computer_id>"`.

### Create a volume (`disk.create_volume`)

#### `-subsystem_id`

The ID of the system on which the volume will be created. The `subsystem_id` is the name of the system as seen in the IBM Tivoli Storage Productivity Center graphical user interface (GUI). This parameter is required.

`-name` The name of the volume. The name is not valid for IBM TotalStorage Enterprise Storage Server and is ignored. If more than one volume is to be created, this value is treated as a prefix.

`-size` The size of the volume.

`-unit` The units to be used while creating the volume. Permissible values are KB, MB, and GB. The default unit is MB.

#### `-number`

The number of volumes to be created.

### SAN Planner (`disk.san_planner`)

#### `-computer_id`

The ID of the computer for which the plan is to be created. The `computer_id` is the name of the computer as seen in the Tivoli Storage Productivity Center GUI.

#### `-subsystem_id`

The ID of the system for which the plan is to be created. The `subsystem_id` is the name of the system as seen in the Tivoli Storage Productivity Center GUI Disk Manager.

#### `-plan_id`

The ID of the previously saved plan. When the `computer_id` and `subsystem_id` are provided, `plan_id` is ignored. When you open a plan in the SAN Planner using launch in context, you must enter the exact name of the plan. The value for `-plan_id` is case-sensitive and must match the name of the plan when it was created. For example, if you created a plan named **UserPlan1**, you must use `-plan_id UserPlan1` in the launch in context command.

### Capacity report (`data.filesystem_capacity_by_computer`)

#### `-computer_id`

The ID of the computer for which the report is to be created. If this value is not specified, the Tivoli Storage Productivity Center GUI opens the default report panel for **Data Manager > Reporting > Capacity > FileSystem Capacity > By Computer**.

You must enter a value for **-computer\_id** that matches the ID for that computer that is stored in the Tivoli Storage Productivity Center database repository. You will receive a warning message if the ID for the computer you enter in the command line does not match the ID stored in the data repository. To ensure that this report is displayed properly when using launch in context, make sure to enter an ID for the computer that matches the ID that is displayed for it through the Topology Viewer or appropriate report in the product's user interface.

### Wasted space report (TPC.reports\_data.wasted\_space)

This report has no parameters. The Tivoli Storage Productivity Center GUI opens the default report panel for **Data Manager > Reporting > System Reports > Data > Wasted Space**.

### Available windows

You can use windows with the inbound launch in context feature to start IBM Tivoli Storage Productivity Center. These windows are called with the **-function window name** argument.

### Available window names for function argument

Table 59. Available windows for inbound launch in context

| Tree Node                                                                                          | Window Name                   |
|----------------------------------------------------------------------------------------------------|-------------------------------|
| Administrative Services->Data Sources->CIMOM Agents                                                | datasource.cimom              |
| Administrative Services->Data Sources->Data Agents                                                 | datasource.data               |
| Administrative Services->Data Sources->Inband fabric Agents                                        | datasource.inband_fabric      |
| Administrative Services->Data Sources->Out of band fabric Agents                                   | datasource.out_of_band_fabric |
| Administrative Services->Data Sources->IBM Tivoli Storage Productivity Center Servers <sup>1</sup> | datasource.tpc_servers        |
| Administrative Services->Data Sources->VMWare VI data source <sup>1</sup>                          | datasource.vmware             |
| IBM Tivoli Storage Productivity Center->Alerting->Alert Log->All                                   | alert.all                     |
| IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Alerts Directed to user               | alert.alerts_directed_to_user |
| IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Storage Subsystem                     | alert.storage_subsystem       |
| IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Computer                              | alert.computer                |
| IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Disk                                  | alert.disk                    |
| IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Filesystem                            | alert.filesystem              |
| IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Directory <sup>1</sup>                | alert.directory               |
| IBM Tivoli Storage Productivity Center->Alerting->Alert Log->User <sup>1</sup>                     | alert.user                    |



Table 59. Available windows for inbound launch in context (continued)

| Tree Node                                                                                                                | Window Name                                   |
|--------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------|
| IBM Tivoli Storage Productivity Center->Alerting->Alert Log->OS User Group <sup>1</sup>                                  | alert.os_user_group                           |
| IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Fabric                                                      | alert.fabric                                  |
| IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Switch                                                      | alert.switch                                  |
| IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Endpoint Device                                             | alert.endpoint_device                         |
| IBM Tivoli Storage Productivity Center->Alerting->Alert Log->External                                                    | alert.external                                |
| IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Tape Library                                                | alert.tape_library                            |
| IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Configuration Analysis <sup>1</sup>                         | alert.configuration_analysis                  |
| IBM Tivoli Storage Productivity Center->Alerting->Alert Log->Hypervisor <sup>1</sup>                                     | alert.hypervisor                              |
| IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data->Disk Space Summary <sup>1</sup>                | data_report.disk_capacity                     |
| IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data->Storage Access Times <sup>1</sup>              | data_report.storage_access_times              |
| IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data->Most Obsolete Files <sup>1</sup>               | data_report.most_obsolete_files               |
| IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data->Storage Availability                           | data_report.storage_availability              |
| IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data->Storage Capacity <sup>1</sup>                  | data_report.storage_capacity                  |
| IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data->Total Freespace <sup>1</sup>                   | data_report.total_freespace                   |
| IBM Tivoli Storage Productivity Center->My Reports->System Reports->Data-> User Space Usage <sup>1</sup>                 | data_report.user_space_usage                  |
| IBM Tivoli Storage Productivity Center->My Reports->System Reports->Disk->Port Performance <sup>1</sup>                  | disk_report.port_performance                  |
| IBM Tivoli Storage Productivity Center->My Reports->System Reports->Disk->Subsystem Performance <sup>1</sup>             | disk_report.subsystem_performance             |
| IBM Tivoli Storage Productivity Center->My Reports->System Reports->Disk->Top Volumes Data Rate Performance <sup>1</sup> | disk_report.top_volumes_data_rate_performance |
| IBM Tivoli Storage Productivity Center->My Reports->System Reports->Disk-> Top Volumes I/O Rate Performance <sup>1</sup> | disk_report.top_volumes_io_rate_performance   |
| IBM Tivoli Storage Productivity Center->My Reports->System Reports->Fabric-> San Assets (ALL)                            | fabric_report.san_assets                      |
| IBM Tivoli Storage Productivity Center->My Reports->System Reports->Fabric ->Port Connections                            | fabric_report.port_connections                |
| IBM Tivoli Storage Productivity Center->My Reports->System Reports->Fabric -> Switch Performance <sup>1</sup>            | fabric_report.switch_performance              |

Table 59. Available windows for inbound launch in context (continued)

| Tree Node                                                                                                                        | Window Name                                          |
|----------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------|
| IBM Tivoli Storage Productivity Center->My Reports->System Reports->Fabric ->Top Switch Ports Data Rate Performance <sup>1</sup> | fabric_report.top_switch_ports_data_rate_performance |
| IBM Tivoli Storage Productivity Center->Topology->Computers                                                                      | topology.computers                                   |
| IBM Tivoli Storage Productivity Center->Topology->Fabrics                                                                        | topology.fabrics                                     |
| IBM Tivoli Storage Productivity Center->Topology->Switches                                                                       | topology.switches                                    |
| IBM Tivoli Storage Productivity Center->Topology->Storage                                                                        | topology.storage                                     |
| IBM Tivoli Storage Productivity Center->Topology->Storage resource groups                                                        | topology.storage_resource_groups                     |
| IBM Tivoli Storage Productivity Center->Topology->Other                                                                          | topology.other                                       |
| IBM Tivoli Storage Productivity Center->Configuration Utility                                                                    | config                                               |
| Disk Manager->Storage Subsystems                                                                                                 | disk.storage_subsystems                              |
| Fabric Manager->Fabrics                                                                                                          | fabric.fabrics                                       |
| Tape Manager->Tape Libraries                                                                                                     | tape.tape_libraries                                  |
| <b>Windows that correspond to launch in context tasks</b>                                                                        |                                                      |
| Fabric Manager -> SAN Planner <sup>1</sup>                                                                                       | disk.san_planner                                     |
| Disk Manager -> Storage Subsystems                                                                                               | disk.create_volume                                   |
| Data Manager-> Reporting -> Capacity -> Filesystem Capacity -> By Computer <sup>1</sup>                                          | data.filesystem_capacity_by_computer                 |
| IBM Tivoli Storage Productivity Center -> My Reports -> System Reports -> Data -> Wasted Space <sup>1</sup>                      | TPC.reports.data.wasted_space                        |
| <sup>1</sup> This window is not available in IBM Tivoli Storage Productivity Center Basic Edition.                               |                                                      |

## Alerts

You can set up Tivoli Storage Productivity Center so that it examines the data that it collects about your storage infrastructure and writes an alert to a log when an event occurs. You also can specify that an action is initiated, such as sending an SNMP trap, sending an email, or running a script when the event occurs.

Alerts that are triggered based on the data collected by monitoring jobs (for example, pings, scans, and probes) must be defined before the monitoring jobs are run. Alerts that are triggered based on an external source, such as an SNMP trap, must be defined after the SNMP server is configured properly.

You can define an alert in the following ways:

- As part of a data collection schedule
- As part of an alert definition

When an event occurs and triggers an alert, the alert is written to a log. You can also select one or more other ways to be notified of the event. These alert

notifications include SNMP traps, IBM Tivoli Enterprise Console events, Tivoli Storage Productivity Center login notifications, operating-system event logs, or email.

**Note:** Alerts are not generated in a Tivoli Storage Productivity Center instance for actions that you perform from that instance. For example, if you start Tivoli Storage Productivity Center and use Disk Manager to assign or unassign volumes for a subsystem, you do not receive alerts for those volume changes in that instance. However, if you assign and unassign volumes outside of that Tivoli Storage Productivity Center instance, an alert is generated.

## Triggering conditions for alerts

In general, the following types of conditions can trigger alerts:

- A data collection schedule failed to complete.
- A change occurred in the storage infrastructure.
- A performance threshold was violated.

The specific conditions that can trigger events vary, depending on the type of storage resource that is being monitored. See “Triggering conditions for alerts” on page 1158 for more information about triggering conditions.

## Alert notifications and triggered actions

Notification methods, or triggered actions, define the method by which you are notified of an alert: SNMP Trap, Tivoli Enterprise Console Event, Login Notification, Windows Event Log, UNIX Syslog, and email. Alerts are always written to the error log. Additionally, if a Data agent or Storage Resource agent is deployed on a monitored computer, you can run a script or start an IBM Tivoli Storage Manager job in response to the alert. See “Triggered actions for alerts” on page 1172 for more information about triggered actions.

## Prerequisites for successfully using alerts

The following conditions must be met in order to successfully use alerts:

- Data collection schedules are configured and scheduled to run on a regular basis.
- If you want to be notified about an alert in some way other than an entry in the log file, such as using SNMP traps, Tivoli Enterprise Console events, or email, alert notification instructions must be configured before using the alert. See “Configuring alert destinations” on page 394).
- If an alert is triggered based on an SNMP trap from the monitored resource, you must properly configure the SNMP server of the monitor resource to enable Tivoli Storage Productivity Center to listen to SNMP traps. The default port number is 162, and the default community is public.

## Default alerts

The following types of alerts are set for your system by default:

- New entity found
- Job failed

You can change the triggered actions and enable or disable these alerts, but you cannot delete them.

## Triggering conditions for alerts

You can define alerts so that IBM Tivoli Storage Productivity Center notifies you when a specified event occurs. Such events are the *triggering conditions* for the alert. The specific triggering conditions that you can use to define an alert depend on the sort of storage resource that you are monitoring.

Some triggering conditions, like thresholds, require you to enter values. If input is not required for the triggering condition you select, the value fields will be grayed out. For information about thresholds, see “Performance thresholds” on page 1176.

### Triggering conditions for computer alerts

Use this topic to view the events that can trigger computer alerts and their related error message identifiers.

Table 60. Triggering conditions and related error messages for computer alerts

| Triggering Condition                     | Explanation                                                                   | Related Error Message* |
|------------------------------------------|-------------------------------------------------------------------------------|------------------------|
| RAM Increased                            | The amount of RAM on a managed computer is increased.                         | ALR0001W               |
| RAM Decreased                            | The amount of RAM on a managed computer is decreased.                         | ALR0001W               |
| Virtual Memory Increased                 | The amount of virtual memory on a managed computer is increased.              | ALR0002W               |
| Virtual Memory Decreased                 | The amount of virtual memory on a managed computer is decreased.              | ALR0002W               |
| New Disk Detected                        | A new disk is discovered on a managed computer.                               | ALR4323W               |
| Disk Not Found                           | A disk is removed from a managed computer.                                    | ALR4324W               |
| New Disk Defect Found                    | A new grown disk defect is detected on a managed disk.                        | ALR0011W               |
| Grown Disk Defects Exceed <i>n</i> value | The grown disk defect threshold is exceeded. Enter a value for <i>n</i> .     | ALR0012W               |
| Disk Failure Predicted                   | A managed disk predicts that a disk failure is imminent.                      | ALR0007W               |
| New Filesystem Detected                  | A new file system is discovered on a managed computer.                        | ALR0005W               |
| Cluster Resource Group Added             | A new cluster resource group is added.                                        | ALR0055W               |
| Cluster Resource Group Removed           | A cluster resource group is removed.                                          | ALR0056W               |
| Cluster Resource Group Moved             | A cluster resource group is moved.                                            | ALR0057W               |
| Computer Unreachable                     | A managed computer cannot be reached.                                         | ALR0017W               |
| Computer Discovered                      | A new unmanaged computer is discovered.                                       | ALR0022W               |
| Computer Status Change Offline           | A managed computer goes offline.                                              | ALR4273W               |
| Computer Status Change Online            | A managed computer comes online.                                              | ALR4274W               |
| Computer Property Change                 | An HBA or HBA to computer node is newly discovered, missing, or rediscovered. | ALR4275W               |
| HBA Driver Version Change                | The driver version of host bus adapter changes.                               | ALR4276W               |

Table 60. Triggering conditions and related error messages for computer alerts (continued)

| Triggering Condition                                                                                                                                                                                                                                                                                                                        | Explanation                                       | Related Error Message* |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------|------------------------|
| HBA Firmware Version Change                                                                                                                                                                                                                                                                                                                 | The firmware version of host bus adapter changes. | ALR4277W               |
| <b>Note:</b> * Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the <i>Messages</i> section of the <i>IBM Tivoli Storage Productivity Center Information Center</i> or the <i>IBM Tivoli Storage Productivity Center Messages Guide</i> . |                                                   |                        |

### Triggering conditions for file system alerts

This topic provides descriptions and related error message identifiers for the triggering conditions that can be set for file system alerts.

Table 61. Triggering conditions and related error messages for file system alerts

| Triggering Condition                         | Explanation                                                                                                                                                                             | Related Error Message* |
|----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
| File System Was Not Found                    | A file system was removed or unmounted from a cluster. This alert applies to a file system on a Storwize V7000 Unified or a host managed by a Storage Resource agent.                   | ALR0006W               |
| File System Was Reconfigured                 | The physical space definition of a managed file system was reconfigured. This alert applies to a file system on a Storwize V7000 Unified or a host managed by a Storage Resource agent. | ALR0008W               |
| File System Free Space Less Than             | A file system is low on free space. This value is usually entered as a percent.                                                                                                         | ALR0009W               |
| Available UNIX File System I-Nodes Less Than | A UNIX file system is low on free i-nodes. This value is usually as a percent.                                                                                                          | ALR0010W               |
| New File System Was Detected                 | A file system was detected for the first time. This alert applies to a file system on a Storwize V7000 Unified or a host managed by a Storage Resource agent.                           | ALR4427I               |
| File System Is Low on Relative Free Space    | A file system is low on free space based on a percentage value threshold. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 Unified.             | ALR4433W               |
| File System Is Low on Absolute Free Space    | A file system is low on free space based on an absolute value threshold. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 Unified.              | ALR4433W               |
| File System Mount State Changed to Normal    | The mount state for a file system reached normal level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 Unified.                               | ALR4358I               |
| File System Mount State Changed to Error     | The mount state for a file system reached error level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 Unified.                                | ALR4356E               |
| File System Is Low On Relative Free I-Nodes  | A file system is low on free i-nodes based on a percentage value threshold. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 Unified.           | ALR4458W               |
| File System Is Low On Absolute Free I-Nodes  | A file system is low on free i-nodes based on an absolute value threshold. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 Unified.            | ALR4458W               |

Table 61. Triggering conditions and related error messages for file system alerts (continued)

| Triggering Condition                                                                                                                                                                                                                                                                                                                        | Explanation                                                                                                                              | Related Error Message* |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
| New File Set Was Detected                                                                                                                                                                                                                                                                                                                   | A file set was detected for the first time. This alert only applies to Storwize V7000 Unified.                                           | ALR4447I               |
| File Set Is Missing                                                                                                                                                                                                                                                                                                                         | A file set that was found during a previous probe was not found in the current probe. This alert only applies to Storwize V7000 Unified. | ALR4474W               |
| File Set Is Linked                                                                                                                                                                                                                                                                                                                          | A file set was linked to a file system. This alert only applies to Storwize V7000 Unified.                                               | ALR4448I               |
| File Set Is Unlinked                                                                                                                                                                                                                                                                                                                        | A file set was unlinked from a file system. This alert only applies to Storwize V7000 Unified.                                           | ALR4455I               |
| File Set Was Rediscovered                                                                                                                                                                                                                                                                                                                   | A file set that was missing was found by the current probe. This alert only applies to Storwize V7000 Unified.                           | ALR4475I               |
| Quota Was Exceeded                                                                                                                                                                                                                                                                                                                          | A quota was exceeded. This alert only applies to Storwize V7000 Unified.                                                                 | ALR4482W               |
| New Quota Was Detected                                                                                                                                                                                                                                                                                                                      | A quota was detected for the first time. This alert only applies to Storwize V7000 Unified.                                              | ALR4496I               |
| <b>Note:</b> * Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the <i>Messages</i> section of the <i>IBM Tivoli Storage Productivity Center Information Center</i> or the <i>IBM Tivoli Storage Productivity Center Messages Guide</i> . |                                                                                                                                          |                        |

### Triggering conditions for directory alerts

Use this topic to view the events that can trigger directory alerts and their related error message identifiers.

Table 62. Triggering conditions and related error messages for directory alerts

| Triggering Condition                                                                                                                                                                                                                                                                                                                        | Explanation                                                                         | Related Error Message* |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|------------------------|
| Directory Not Found                                                                                                                                                                                                                                                                                                                         | A directory cannot be found.                                                        | ALR0014W               |
| Directory Consumes More Than                                                                                                                                                                                                                                                                                                                | A user or directory storage quota is exceeded. Enter a value, usually in megabytes. | ALR0015W               |
| <b>Note:</b> * Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the <i>Messages</i> section of the <i>IBM Tivoli Storage Productivity Center Information Center</i> or the <i>IBM Tivoli Storage Productivity Center Messages Guide</i> . |                                                                                     |                        |

### Triggering conditions for instance alerts

Use this topic lists to view the events that can trigger instance alerts and their related error message identifiers.

Table 63. Triggering conditions and related error messages for instance alerts

| Database Application | Triggering Condition                     | Explanation                                                           | Related Error Message* |
|----------------------|------------------------------------------|-----------------------------------------------------------------------|------------------------|
| All                  | New Database - Tablespace Discovered     | A new database-tablespace is discovered.                              | ALR0028W               |
| Oracle               | Archive Log Directory Contains More Than | An archived log directory quota is exceeded. Enter a value for quota. | ALR0027W               |
|                      | New Tablespace Discovered                | A new tablespace is discovered.                                       | ALR0028W               |
| SQL Server           | New Database Discovered                  | A new database is discovered.                                         | ALR0028W               |

Table 63. Triggering conditions and related error messages for instance alerts (continued)

| Database Application                                                                                                                                                                                                                                                                                                                              | Triggering Condition                | Explanation                                                               | Related Error Message* |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|---------------------------------------------------------------------------|------------------------|
| Sybase                                                                                                                                                                                                                                                                                                                                            | Total Device Freespace Less Than    | The total amount of device free space is less than a specified amount.    | ALR0044W               |
|                                                                                                                                                                                                                                                                                                                                                   | Total Device Freespace Greater Than | The total amount of device free space is greater than a specified amount. | ALR0045W               |
|                                                                                                                                                                                                                                                                                                                                                   | New Device Discovered               | A new device is discovered.                                               | ALR0042W               |
|                                                                                                                                                                                                                                                                                                                                                   | Device Dropped                      | A device is dropped.                                                      | ALR0043W               |
|                                                                                                                                                                                                                                                                                                                                                   | New Database Discovered             | A new database is discovered.                                             | ALR0028W               |
| DB2                                                                                                                                                                                                                                                                                                                                               | New Tablespace Discovered           | A new database-tablespace is discovered.                                  | ALR0028W               |
| <p><b>Note:</b> * Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the <i>Messages</i> section of the <i>IBM Tivoli Storage Productivity Center Information Center</i> or the <i>IBM Tivoli Storage Productivity Center Messages Guide</i>.</p> |                                     |                                                                           |                        |

### Triggering conditions for database-tablespace alerts

Use this topic to view the events that can trigger database-tablespace alerts and their related error message identifiers.

Table 64. Triggering conditions and related error messages for database-tablespace alerts

| Database application | Triggering Condition                      | Explanation                                                                                  | Related Error Message* |
|----------------------|-------------------------------------------|----------------------------------------------------------------------------------------------|------------------------|
| All                  | Database - Tablespace Freespace Less Than | The free space on a database-tablespace is less than a specified amount.                     | ALR0031W               |
|                      | Database - Tablespace Offline             | A database-tablespace is taken offline.                                                      | ALR0030W               |
|                      | Database - Tablespace Dropped             | A database-tablespace is dropped.                                                            | ALR0029W               |
| Oracle               | Tablespace Freespace Less Than            | The free space on a tablespace is less than a specified amount.                              | ALR0031W               |
|                      | Freespace Fragmented into More Than       | The free space on a tablespace is fragmented into more than the specified number of extents. | ALR0032W               |
|                      | Largest Free Extent Less Than             | The largest free extent is less than the specified size.                                     | ALR0033W               |
|                      | Tablespace Offline                        | A tablespace is taken offline                                                                | ALR0030W               |
|                      | Tablespace Dropped                        | A tablespace is dropped.                                                                     | ALR0029W               |
| SQL Server           | Database Data Freespace Less Than         | The free space on a database is less than a specified amount.                                | ALR0031W               |
|                      | Database Log Freespace Less Than          | The free space on a database log is less than the specified percent.                         | ALR0041W               |
|                      | Database Offline                          | A database is taken offline.                                                                 | ALR0030W               |
|                      | Database Dropped                          | A database is dropped.                                                                       | ALR0029W               |

Table 64. Triggering conditions and related error messages for database-tablespace alerts (continued)

| Database application                                                                                                                                                                                                                                                                                                                              | Triggering Condition              | Explanation                                                          | Related Error Message* |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------|----------------------------------------------------------------------|------------------------|
| Sybase                                                                                                                                                                                                                                                                                                                                            | Database Data Freespace less Than | The free space on a database is less than a specified amount.        | ALR0031W               |
|                                                                                                                                                                                                                                                                                                                                                   | Database Log Freespace Less Than  | The free space on a database log is less than the specified percent. | ALR0027W               |
|                                                                                                                                                                                                                                                                                                                                                   | Last Dump Time Less Than          | The last dump time was the less than the specified number of days.   | ALR1111W               |
|                                                                                                                                                                                                                                                                                                                                                   | Database Offline                  | A database is taken offline.                                         | ALR0030W               |
|                                                                                                                                                                                                                                                                                                                                                   | Database Dropped                  | A database is dropped                                                | ALR0029W               |
| DB2                                                                                                                                                                                                                                                                                                                                               | Tablespace Freespace Less Than    | The free space on a tablespace is less than a specified amount.      | ALR0031W               |
|                                                                                                                                                                                                                                                                                                                                                   | Tablespace Offline                | A tablespace is taken offline.                                       | ALR0030W               |
|                                                                                                                                                                                                                                                                                                                                                   | Tablespace Dropped                | A tablespace is dropped                                              | ALR0029W               |
| <p><b>Note:</b> * Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the <i>Messages</i> section of the <i>IBM Tivoli Storage Productivity Center Information Center</i> or the <i>IBM Tivoli Storage Productivity Center Messages Guide</i>.</p> |                                   |                                                                      |                        |

### Triggering conditions for table alerts

Use this topic to view the events that can trigger table alerts and their related error message identifiers.

Table 65. Triggering conditions and related error messages for table alerts

| Database application | Triggering Condition             | Explanation                                                  | Related Error Message* |
|----------------------|----------------------------------|--------------------------------------------------------------|------------------------|
| All                  | Total Table Size Greater Than    | The total size of a table is greater than a specified size.  | ALR0036W               |
|                      | Table Dropped                    | A table is dropped.                                          | ALR0039W               |
| Oracle               | Max Extents - Allocated Extents  | A table segment is near its maximum allowed extents.         | ALR0035W               |
|                      | Segment Has More Than            | A table segment has taken too many extents.                  | ALR0034W               |
|                      | Chained Row Count Greater Than   | A table has too many chained rows.                           | ALR0037W               |
|                      | Empty Used Segment Space Exceeds | An empty used segment space exceeds the specified size.      | ALR0038W               |
|                      | Total Table Size Greater Than    | The total size of a table is greater than a specified size.  | ALR0036W               |
|                      | Table Dropped                    | A table is dropped.                                          | ALR0039W               |
| SQL Server           | Total Table Size Greater Than    | The total size of a table is greater than a specified size.  | ALR0036W               |
|                      | Table Dropped                    | A table is dropped.                                          | ALR0039W               |
| Sybase               | Forwarded Row Count Greater Than | A forwarded row count is greater than the specified percent. | ALR0037W               |
|                      | Total Table Size Greater Than    | The total size of a table is greater than a specified size.  | ALR0036W               |
|                      | Table Dropped                    | A table is dropped.                                          | ALR0039W               |



Table 65. Triggering conditions and related error messages for table alerts (continued)

| Database application | Triggering Condition            | Explanation                                                  | Related Error Message* |
|----------------------|---------------------------------|--------------------------------------------------------------|------------------------|
| DB2                  | Overflow Row Count Greater Than | An overflow row count is greater than the specified percent. |                        |
|                      | Total Table Size Greater Than   | The total size of a table is greater than a specified size.  | ALR0036W               |
|                      | Table Dropped                   | A table is dropped.                                          | ALR0039W               |

**Note:** \* Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the *Messages* section of the *IBM Tivoli Storage Productivity Center Information Center* or the *IBM Tivoli Storage Productivity Center Messages Guide*.

### Triggering conditions for storage subsystem alerts

Use this topic to view the events that can trigger storage subsystem alerts and their related error message identifiers.

Table 66. Triggering conditions and related error messages for storage-subsystem alerts

| Triggering Condition                     | Explanation                                                          | Related Error Message <sup>1</sup> |
|------------------------------------------|----------------------------------------------------------------------|------------------------------------|
| New Storage Subsystem discovered         | A new storage subsystem is discovered.                               | ALR1114M                           |
| Cache Increased                          | A subsystem cache increases.                                         | ALR0050W                           |
| Cache Decreased                          | A subsystem cache decreases.                                         | ALR0050W                           |
| New Disk Detected <sup>2</sup>           | A new disk is detected.                                              | ALR4323W                           |
| Disk Not Found <sup>2</sup>              | A previously detected disk is not found.                             | ALR4324W                           |
| Storage Subsystem Not Found              | A previously detected storage subsystem is not found.                | ALR1113W                           |
| Subsystem Status Change Offline          | A subsystem goes offline.                                            | ALR4241W                           |
| Subsystem Status Change Online           | A subsystem comes online.                                            | ALR4242I                           |
| Subsystem Version Change                 | A subsystem version changes.                                         | ALR4243W                           |
| Subsystem Property Change                | A port on a subsystem is newly discovered, missing, or rediscovered. | ALR4278W                           |
| Subsystem Allocated Capacity Change      | The allocated capacity of a subsystem changes.                       | ALR4244W                           |
| Subsystem Available Capacity Change      | The available capacity of a subsystem changes.                       | ALR4245W                           |
| Subsystem Backend Capacity Change        | The allocated capacity of a subsystem changes.                       | ALR4246W                           |
| Backend Controller Status Change Offline | A back-end controller goes offline.                                  | ALR4247W                           |
| Backend Controller Status Change Online  | A back-end controller comes online.                                  | ALR4248I                           |
| Volume Status Change Offline             | A volume goes offline.                                               | ALR4249W                           |
| Volume Status Change Online              | A volume comes online.                                               | ALR4250I                           |

Table 66. Triggering conditions and related error messages for storage-subsystem alerts (continued)

| Triggering Condition                 | Explanation                                                                                                                                                                                                                          | Related Error Message <sup>1</sup>     |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------|
| Volume Capacity Change               | The capacity of a volume changes.                                                                                                                                                                                                    | ALR4251W                               |
| Pool Status Change                   | A pool is missing or rediscovered.                                                                                                                                                                                                   | ALR4252W                               |
| Pool Discovered                      | A pool is discovered.                                                                                                                                                                                                                | ALR4253I                               |
| Pool Status Change Offline           | A pool goes offline.                                                                                                                                                                                                                 | ALR4254W                               |
| Pool Status Change Online            | A pool comes online.                                                                                                                                                                                                                 | ALR4255I                               |
| Pool Capacity Change                 | The capacity of a pool changes.                                                                                                                                                                                                      | ALR4256W                               |
| Pool Available Space Change          | The available space in a pool changes.                                                                                                                                                                                               | ALR4257W                               |
| Node Status Change Offline           | A node goes offline.                                                                                                                                                                                                                 | ALR4226W                               |
| Node Status Change Online            | A node comes online.                                                                                                                                                                                                                 | ALR4227W                               |
| Node State Change                    | A IBM System Storage SAN Volume Controller node is missing or rediscovered.                                                                                                                                                          | ALR4224W                               |
| Node Discovered                      | A SAN Volume Controller node is discovered.                                                                                                                                                                                          | ALR4225W                               |
| Datapath State Change                | The state of a data path changes.                                                                                                                                                                                                    | ALR4283W                               |
| Datapath Discovered                  | A data path is discovered.                                                                                                                                                                                                           | ALR4280I                               |
| New Volume discovered                | A new volume is discovered during a storage subsystem probe.                                                                                                                                                                         | ALR4325W                               |
| Volume not found                     | A volume which had existed before is not found during a storage subsystem probe.                                                                                                                                                     | ALR4326W                               |
| Link Failure Rate Threshold          | The average number of miscellaneous fibre channel link errors (for example, an unexpected Not Operational (NOS) received or a link state machine failure detected) per second that were encountered by a port over a period of time. |                                        |
| Error Frame Rate Threshold           | The threshold for the number of frames per second that violated Fibre Channel protocol for a port over a period of time.                                                                                                             |                                        |
| Overall Port Response Time Threshold | The threshold for the average number of milliseconds that it took to service each operation (send and receive), by a particular port over a particular time interval.                                                                | ALR0540W, ALR0541W, ALR0542W, ALR0543W |
| Total Port Data Rate Threshold       | The threshold for the average number of megabytes transferred (send and receive) for a particular port over a particular time interval                                                                                               | ALR0536W, ALR0537W, ALR0538W, ALR0539W |
| Total Port IO Rate Threshold         | The threshold for the average number of send and receive operations processed per second, for a particular port over a particular time interval                                                                                      | ALR0532W, ALR0533W, ALR0534W, ALR0535W |
| CPU Utilization Threshold            | The threshold is based on the average utilization percentage of the CPUs in the SAN Volume Controller nodes. This value for each node is checked against the threshold boundaries for each collection interval.                      | ALR0556W, ALR0557W, ALR0558W, ALR0559W |

Table 66. Triggering conditions and related error messages for storage-subsystem alerts (continued)

| Triggering Condition                                     | Explanation                                                                                                                                                                                                             | Related Error Message <sup>1</sup>      |
|----------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------|
| Cache Holding Time Threshold                             | The threshold is reached for the average cache holding time, in seconds, for I/O data in a particular component over a particular time interval. Shorter time periods indicate adverse performance                      | ALR0528W, ALR0529W, ALR0530W, ALR0531W  |
| Write-cache Delay Percentage Threshold                   | The threshold is based on the percentage of I/O operations that were delayed due to write-cache space constraints. This value for each node is checked against the threshold boundaries for each collection interval    | ALR0524W, ALR0525W, ALR0526W, ALR0527W  |
| Total Data Rate Threshold                                | For the storage-subsystem controllers or I/O groups, the threshold for the average number of MB per second for read and write operations is reached                                                                     | ALR0520W, ALR0521W, ALR0522W, ALR0523W  |
| Total I/O Rate Threshold                                 | For the storage subsystem controllers or I/O groups, the threshold for the average number of I/O operations per second is reached                                                                                       | ALR0516W, ALR0517W, ALR0518W, ALR0519W  |
| Overall Backend Response Time Threshold                  | The threshold for the average number of milliseconds that it takes to respond to MDisk I/O operations is reached                                                                                                        | ALR0512W, ALR0513W, ALR0514W, ALR0515W  |
| Total Backend Data Rate Threshold                        | For MDisk groups, the threshold for the average number of megabytes per second that are transferred for MDisk read and write operations is reached                                                                      | ALR0508W, ALR0509W, ALR0510W, ALR0511W  |
| Zero Buffer-Buffer Credit Timer                          | The number of microseconds for which the port has been unable to send frames due to lack of buffer credit since the last node reset.                                                                                    |                                         |
| Total Backend I/O Rate Threshold                         | For MDisk groups, the threshold for the average number of MDisk read and write operations per second is reached.                                                                                                        | ALR0504W , ALR0505W, ALR0506W, ALR0507W |
| Invalid Transmission Word Rate Threshold                 | The threshold for the average number of times per second that "bit" errors are detected on a port over a period of time.                                                                                                |                                         |
| CRC Error Rate Threshold                                 | The threshold for the average number of frames received per second where the CRC in the frame did not match the CRC computed by the receiver for a component over a period of time.                                     |                                         |
| Disk Utilization Percentage Threshold                    | The threshold is reached for the approximate utilization percentage of a particular array over a particular time interval. For example, the average percentage of time that the disk associated with the array was busy | ALR0500W, ALR0501W, ALR0502W, ALR0503W  |
| Port Receive Bandwidth Percentage Threshold <sup>3</sup> | The threshold for the average port bandwidth utilization percentage for receive operations.                                                                                                                             |                                         |
| Port Send Bandwidth Percentage Threshold <sup>3</sup>    | The threshold for the average port bandwidth utilization percentage for send operations.                                                                                                                                |                                         |
| Port Receive Utilization Percentage Threshold            | The threshold for the average amount of time that ports are busy receiving data.                                                                                                                                        |                                         |
| Port Send Utilization Percentage Threshold               | The threshold for the average port bandwidth utilization percentage for send operations.                                                                                                                                |                                         |
| Peak Backend Write Response Time Threshold               | The threshold for the peak (worst) response time among all MDisk write operations by a node.                                                                                                                            |                                         |

Table 66. Triggering conditions and related error messages for storage-subsystem alerts (continued)

| Triggering Condition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Explanation                                                                                                                                                              | Related Error Message <sup>1</sup> |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------|
| Non-preferred Node Usage Percentage Threshold                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | The threshold for the non-preferred node usage percentage of an I/O group.                                                                                               |                                    |
| Port to Local Node Receive Queue Time Threshold                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | The threshold for the average number of milliseconds that each receive operation issued to another node in the local SVC cluster spent on the queue before being issued. |                                    |
| Port to Local Node Send Queue Time Threshold                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | The threshold for the average number of milliseconds that each send operation issued to another node in the local SVC cluster spent on the queue before being issued.    |                                    |
| Port to Local Node Receive Response Time Threshold                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | The threshold for the average number of milliseconds it took to service each receive operation from another node in the local SVC cluster.                               |                                    |
| Port to Local Node Send Response Time Threshold                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | The threshold for the average number of milliseconds it took to service each send operation from another node in the local SVC cluster.                                  |                                    |
| Backend Write Queue Time Threshold                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | The threshold for the average number of milliseconds that each write operation spent on the queue before being issued to the back-end device.                            |                                    |
| Backend Read Queue Time Threshold                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | The threshold for the average number of milliseconds that each read operation spent on the queue before being issued to the back-end device.                             |                                    |
| Backend Write Response Time Threshold                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | The threshold for the average number of milliseconds that it took to service each array and MDisk write operation.                                                       |                                    |
| Backend Read Response Time Threshold                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | The threshold for the average number of milliseconds that it took to service each array and MDisk read operation.                                                        |                                    |
| <p><b>Notes:</b></p> <ol style="list-style-type: none"> <li>1. Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the <i>Messages</i> section of the <i>IBM Tivoli Storage Productivity Center Information Center</i> or the <i>IBM Tivoli Storage Productivity Center Messages Guide</i>.</li> <li>2. This condition does not trigger alerts for the following storage systems: <ul style="list-style-type: none"> <li>• SAN Volume Controller</li> <li>• Storwize V7000</li> <li>• Storwize V7000 Unified</li> </ul> </li> <li>3. This condition does not trigger alerts for the DS6000.</li> </ol> |                                                                                                                                                                          |                                    |

**Related tasks:**

“Creating a performance threshold within an alert” on page 150

You can monitor the performance of your enterprise by creating alerts with performance thresholds for storage subsystems and SAN switches. By defining alerts that are triggered by performance thresholds, you can get timely information about performance issues in your enterprise.

**Related reference:**

“Performance thresholds” on page 1176

Performance thresholds are triggering conditions which are used to monitor a component with user-defined values.

**Triggering conditions for fabric alerts**

Use this topic to view the events that can trigger fabric alerts and their related error message identifiers.

Table 67. Triggering conditions and related error messages for fabric alerts

| Triggering Condition        | Explanation                                                                      | Related Error Message* |
|-----------------------------|----------------------------------------------------------------------------------|------------------------|
| Fabric Discovered           | A fabric is discovered.                                                          | ALR4236W               |
| Fabric State Change         | A fabric is missing or is rediscovered.                                          | ALR4199W               |
| Fabric to Switch Change     | A switch to fabric association is discovered, rediscovered, or missing.          | ALR4206W               |
| Zone State Change           | A zone is missing or is rediscovered.                                            | ALR4207W               |
| Zone Discovered             | A zone is discovered.                                                            | ALR4237W               |
| Zone Set State Change       | A zone set in a fabric was activated, deactivated, or has gone missing.          | ALR4208W               |
| Zone Set Discovered         | A zone set is discovered.                                                        | ALR4238W               |
| Zone to Zone Alias Change   | A zone to zone alias association is discovered, missing, or rediscovered.        | ALR4209W               |
| Zone to Zone Member Change  | A zone to zone member association is discovered, is missing, or is rediscovered. | ALR4210W               |
| Zone Set to Zone Change     | A zone set to zone association is discovered, is missing, or is rediscovered.    | ALR4211W               |
| New Connection              | A new connection is detected.                                                    | ALR4197W               |
| Connection State Change     | A connection is missing or rediscovered.                                         | ALR4198W               |
| Zone Alias to Member Change | A zone alias to member association is discovered, rediscovered, or is missing.   | ALR4327W               |

**Note:** \* Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the *Messages* section of the *IBM Tivoli Storage Productivity Center Information Center* or the *IBM Tivoli Storage Productivity Center Messages Guide*.

**Triggering conditions for switch alerts**

Use this topic to view the events that can trigger switch alerts and their related error message identifiers.

**Note:** The following triggering conditions are performance thresholds. They trigger alerts only if performance monitors are run on the switch:

- Total Port Data Rate Threshold
- Port Receive Bandwidth Percentage Threshold
- Port Send Bandwidth Percentage Threshold

- Link Failure Rate Threshold
- Error Frame Rate Threshold
- Total Port Packet Rate Threshold

Table 68. Triggering conditions and related error messages for switch alerts

| Triggering Condition                                                                                                                                                                                                                                                                                                                              | Explanation                                                                                                                                                                                                                                                                     | Related Error Message*                 |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------|
| Switch Discovered                                                                                                                                                                                                                                                                                                                                 | A switch is discovered.                                                                                                                                                                                                                                                         | ALR4240W                               |
| Switch State Change                                                                                                                                                                                                                                                                                                                               | A switch is missing or is rediscovered.                                                                                                                                                                                                                                         | ALR4229W                               |
| Switch Property Changes                                                                                                                                                                                                                                                                                                                           | A switch port is discovered, is missing, or is rediscovered.                                                                                                                                                                                                                    | ALR4230W                               |
| Switch Status Degraded                                                                                                                                                                                                                                                                                                                            | The status of the switch has changed from OK to a lower status. This does not necessarily mean that the switch has gone missing.                                                                                                                                                | ALR4213W                               |
| Switch Status Improved                                                                                                                                                                                                                                                                                                                            | The status of the switch has been upgraded to OK.                                                                                                                                                                                                                               | ALR4214W                               |
| Switch Version Change                                                                                                                                                                                                                                                                                                                             | The version of a switch changes.                                                                                                                                                                                                                                                | ALR4215W                               |
| Switch to Port Change                                                                                                                                                                                                                                                                                                                             | A switch to switch port association is discovered, is missing, or is rediscovered.                                                                                                                                                                                              | ALR4216W                               |
| Switch Blade Change                                                                                                                                                                                                                                                                                                                               | A switch module is discovered, is missing, or is rediscovered.                                                                                                                                                                                                                  | ALR4217W                               |
| Switch Blade Change Offline                                                                                                                                                                                                                                                                                                                       | A switch module goes offline.                                                                                                                                                                                                                                                   | ALR4218W                               |
| Switch Blade Change Online                                                                                                                                                                                                                                                                                                                        | A switch module comes online.                                                                                                                                                                                                                                                   | ALR4219W                               |
| Total Port Data Rate Threshold                                                                                                                                                                                                                                                                                                                    | A total port data rate threshold is reached.                                                                                                                                                                                                                                    | ALR0536W, ALR0537W, ALR0538W, ALR0539W |
| Port Receive Bandwidth Percentage Threshold                                                                                                                                                                                                                                                                                                       | The threshold on the average port bandwidth utilization percentage for receive operations. The Port Send Utilization Percentage metric is checked against the threshold boundaries for each collection interval. This threshold uses the following default values: 85,75,-1,-1. |                                        |
| Port Send Bandwidth Percentage Threshold                                                                                                                                                                                                                                                                                                          | The threshold on the average port bandwidth utilization percentage for send operations. The Port Send Utilization Percentage metric is checked against the threshold boundaries for each collection interval. This threshold uses the following default values: 85,75,-1,-1.    |                                        |
| Link Failure Rate Threshold                                                                                                                                                                                                                                                                                                                       | A link failure rate threshold is reached.                                                                                                                                                                                                                                       | ALR0548W, ALR0549W, ALR0550W, ALR0551W |
| Error Frame Rate Threshold                                                                                                                                                                                                                                                                                                                        | An error frame rate threshold is reached.                                                                                                                                                                                                                                       | ALR0544W, ALR0545W, ALR0546W, ALR0547W |
| Total Port Packet Rate Threshold                                                                                                                                                                                                                                                                                                                  | A total port packet rate threshold is reached.                                                                                                                                                                                                                                  | ALR0552W, ALR0553W, ALR0554W, ALR0555W |
| <p><b>Note:</b> * Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the <i>Messages</i> section of the <i>IBM Tivoli Storage Productivity Center Information Center</i> or the <i>IBM Tivoli Storage Productivity Center Messages Guide</i>.</p> |                                                                                                                                                                                                                                                                                 |                                        |

## Triggering conditions for endpoint device alerts

Use this topic to view the events that can trigger endpoint device alerts and their related error message identifiers.

Table 69. Triggering conditions and related error messages for endpoint device alerts

| Triggering Condition                                                                                                                                                                                                                                                                                                                        | Explanation                                       | Related Error Message* |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------|------------------------|
| Endpoint Discovered                                                                                                                                                                                                                                                                                                                         | An endpoint device is discovered.                 | ALR4239W               |
| Endpoint State Change                                                                                                                                                                                                                                                                                                                       | An endpoint device is missing or is rediscovered. | ALR4231W               |
| Endpoint Version Change                                                                                                                                                                                                                                                                                                                     | The version for an endpoint device changes.       | ALR4313W               |
| <b>Note:</b> * Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the <i>Messages</i> section of the <i>IBM Tivoli Storage Productivity Center Information Center</i> or the <i>IBM Tivoli Storage Productivity Center Messages Guide</i> . |                                                   |                        |

### Triggering conditions for hypervisor alerts

Use this topic to view the events that can trigger hypervisor alerts and their related error message identifiers.

Table 70. Triggering conditions and related error messages for hypervisor alerts

| Triggering Condition                                                                                                                                                                                                                                                                                                                        | Explanation                                         | Related Error Message* |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------|------------------------|
| Hypervisor Discovered                                                                                                                                                                                                                                                                                                                       | A hypervisor is discovered.                         | ALR4315W               |
| Hypervisor Missing                                                                                                                                                                                                                                                                                                                          | A hypervisor that was previously probed is missing. | ALR4321W               |
| Virtual Machine Added                                                                                                                                                                                                                                                                                                                       | A new virtual machine was added.                    | ALR4316W               |
| Virtual Machine Removed                                                                                                                                                                                                                                                                                                                     | A virtual machine was removed.                      | ALR4317W               |
| <b>Note:</b> * Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the <i>Messages</i> section of the <i>IBM Tivoli Storage Productivity Center Information Center</i> or the <i>IBM Tivoli Storage Productivity Center Messages Guide</i> . |                                                     |                        |

### Triggering conditions for Other NAS alerts

This topic provides descriptions and related error message identifiers for the triggering conditions that can be set for network-attached storage (NAS) alerts.

Table 71. Triggering conditions and related error messages for NAS alerts

| Triggering Condition              | Explanation                                                                                                                                                                        | Related Error Message* |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
| Node Status Changed to Normal     | The clustered Common Internet File System (CIFS) for a NAS device is active. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module. | ALR4365I               |
| Node Status Changed to Warning    | The clustered CIFS for a NAS device is disabled. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.                             | ALR4366W               |
| Node Status Changed to Error      | The status of a clustered CIFS for a NAS device reached the error level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.     | ALR4367E               |
| GPFS Status Reached Normal Level  | The General Parallel File System (GPFS) for a NAS device is active. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.          | ALR4368I               |
| GPFS Status Reached Warning Level | The GPFS state for a NAS device reached the warning level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.                   | ALR4369W               |

Table 71. Triggering conditions and related error messages for NAS alerts (continued)

| Triggering Condition                  | Explanation                                                                                                                                                                               | Related Error Message* |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
| GPFS Status Reached Error Level       | The GPFS for a NAS device is down. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.                                                  | ALR4370E               |
| Node State Changed                    | The state of a node changed.                                                                                                                                                              | ALR1245W               |
| Node Was Discovered                   | A node was discovered.                                                                                                                                                                    | ALR1246W               |
| Node Status Changed to Online         | A node status changed to online.                                                                                                                                                          | ALR4227W               |
| Node Status Changed to Offline        | A node status changed to offline.                                                                                                                                                         | ALR4226W               |
| Processor Usage Reached Normal Level  | Processor use for a network-attached storage (NAS) device reached the normal level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module. | ALR4361I               |
| Processor Usage Reached Warning Level | Processor use for a NAS device reached the warning level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.                           | ALR4360W               |
| Processor Usage Reached Error Level   | Processor use for a NAS device reached the error level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.                             | ALR4359E               |
| Memory Usage Reached Normal Level     | Memory use for a NAS device reached the normal level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.                               | ALR4364I               |
| Memory Usage Reached Error Level      | Memory use for a NAS device reached the error level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.                                | ALR4362E               |
| Cluster Was Discovered                | A cluster was discovered.                                                                                                                                                                 | ALR4528I               |
| Cluster Was Removed                   | A cluster that was found during a probe was not found in the current probe.                                                                                                               | ALR4529I               |
| Cluster Was Rediscovered              | A cluster that was missing was found by the current probe.                                                                                                                                | ALR4530I               |
| New Export Was Detected               | An export was detected for the first time.                                                                                                                                                | ALR4460I               |
| Export Is Missing                     | An export that was found during a previous probe was not found in the current probe.                                                                                                      | ALR4470W               |
| Export Was Rediscovered               | An export that was missing was found by the current probe.                                                                                                                                | ALR4471I               |
| Export State Changed                  | The state of an export changed.                                                                                                                                                           | ALR4461W               |
| Export Path Changed                   | An export was linked to a new path.                                                                                                                                                       | ALR4462W               |
| Export Protocol Changed               | The protocols that were configured for an export changed.                                                                                                                                 | ALR4463W               |
| New NSD Detected                      | A Network Shared Disk (NSD) was detected for the first time.                                                                                                                              | ALR4503I               |
| NSD Is Missing                        | An NSD that was found during a previous probe was not found in the current probe.                                                                                                         | ALR4505W               |
| NSD State Changed                     | The state of an NSD changed.                                                                                                                                                              | ALR4507W               |



Table 71. Triggering conditions and related error messages for NAS alerts (continued)

| Triggering Condition                                                                                                                                                                                                                                                                                                                              | Explanation                                                                                                                                               | Related Error Message* |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
| NSD Status Reached Normal Level                                                                                                                                                                                                                                                                                                                   | NSD use for a NAS device reached the normal level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.  | ALR4387I               |
| NSD Status Reached Warning Level                                                                                                                                                                                                                                                                                                                  | NSD use for a NAS device reached the warning level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module. | ALR4386W               |
| NSD Status Reached Error Level                                                                                                                                                                                                                                                                                                                    | NSD use for a NAS device reached the error level. This alert is triggered when the corresponding SNMP trap is received from Storwize V7000 File Module.   | ALR4385E               |
| <p><b>Note:</b> * Use the error message identifiers in this column to locate any additional information about the error messages related to triggering conditions in the <i>Messages</i> section of the <i>IBM Tivoli Storage Productivity Center Information Center</i> or the <i>IBM Tivoli Storage Productivity Center Messages Guide</i>.</p> |                                                                                                                                                           |                        |

## Alert suppression actions

If you selected a threshold as a triggering condition for alerts, you can specify the conditions under which alerts are triggered and choose whether you want to suppress repeating alerts. Alerts can be suppressed to avoid generating too many alert log entries or too many actions when the triggering condition occurs often. You can view suppressed alerts in the constraint violation reports.

You can define the following options that enable you to specify conditions that trigger and suppress alerts. If a threshold is not selected as a triggering condition, these options are not available.

### Trigger alerts for both critical and warning conditions

Generates an alert upon the violation of either critical or warning threshold boundaries. This is the default.

### Trigger alerts for critical conditions only

Generates alerts only upon violation of one of the critical threshold boundaries. Violation of a warning boundary creates an entry in the constraint violation report, but does not result in an entry in the alert log or an action being triggered.

### Trigger no alerts

Does not generate an alert upon violation of any threshold boundaries. Creates entries only in the constraint violation report.

### Do not suppress repeating alerts

Does not suppress any repeating alerts. This is the default.

### Suppress alerts unless the Triggering Condition has been violated continuously for a specified length of time

Generates alerts only if the triggering condition has occurred continuously within the length of time specified in the **Length of time** field. Alerts for the first and any subsequent occurrences of the triggering condition within the specified time in minutes will be suppressed. At the point that there have been consecutive occurrences with the specified number of minutes, an alert is generated. When the specified suppression period has expired, the cycle starts again. Note that the timing for this feature is based on the IBM Tivoli Storage Productivity Center server clock rather than the various system clocks.

This option is useful for cases when a single occurrence of the triggering condition may be insignificant, but repeated occurrences can signal a potential problem.

#### **Suppress alerts if a repeat violation has occurred within a specified length of time after the initial violation of the Triggering Condition**

Generates alerts only for the first occurrence of the triggering condition. Alerts for repeated occurrences of the triggering condition within the length of time specified in the **Length of time** field are suppressed. When the specified suppression period has expired, the cycle starts again. Note that the timing for this feature is based on the IBM Tivoli Storage Productivity Center server clock rather than the various system clocks.

This option is useful for avoiding email messages or similar disruptive alerts if the same triggering condition occurs repeatedly in successive sample passes. This option generally useful for all threshold types.

## **Triggered actions for alerts**

You define triggered actions to occur as a result of the alert.

An alert is always entered in the error log file each time it occurs. In addition, you can define the following notifications or triggered actions:

#### **SNMP Trap**

Generates an SNMP trap message to any network-management system, console, or terminal. You must set up the SNMP trap ringer with the `fabric.mib` and `tivoliSRM.mib` files.

#### **TEC / OMNIBus Event**

Sends an alert to the IBM Tivoli Enterprise Console or OMNIBus EIF probe server in your environment. The Tivoli Enterprise Console administrator can write correlation and automation rules to analyze the IBM Tivoli Storage Productivity Center events. Tivoli Enterprise Console also performs responses such as sending further notifications, creating or updating trouble tickets, and running programs. You must load the `tivoliSRM.baroc` and `fabric.baroc` files into the active rulebase of the Tivoli Enterprise Console server.

#### **Login Notification**

Tivoli Storage Productivity Center lists the alerts for the specified users when they log in. Use the **Preferences > Edit General** option on the Tivoli Storage Productivity Center menu bar to control which alerts a user views upon logging in to the product.

#### **Windows Event Log, UNIX Syslog**

Writes the alerts to the operating-system log. You can select from the following levels of operating-system events:

- UNIX: Informational, Notice, Warning, Error, Critical, or Alert
- Microsoft Windows: Error, Warning, and Information.

For UNIX, you also can select User or Local.

#### **Run Script**

Runs a script in response to the alert. Use this option to run third-party tools for such actions as recovery and provisioning.

1. Click **Define** to specify the script that is run when the triggering condition occurs. The **Specify Script** window opens.
2. Type the name of the script in the **Script Name** field.

3. Select where to run the script.

If parameters are given for the triggering condition, they are listed in the **Specify Script** window. These parameters are used when the alert is triggered.

**Note:** Scripts can only be run against systems on which a Storage Resource agent or Data agent is deployed.

#### **Archive/Backup**

(Constraints only) Runs an IBM System Storage Productivity Center backup or archive job in response to an alert associated with a constraint. Click the Define button.

**Email** Sends an email. Use this section to view, add, and delete email addresses and edit the text of the message in the email.

For **Directory Alerts**, if the **Directory Consumes More Than** triggering condition is selected, additional fields are available to send email to the quota violator. The email address can be based on **Directory Owner** or **Directory Name**.

## **Alert display colors**

Status symbols for alerts are shown in the colors listed for the alert action.

When you have active alerts shown within the subnodes under the **IBM Tivoli Storage Productivity Center-> Alerting -> Alert Log** node, the indicator next to those subnodes will appear as green or red depending on their corresponding alerts. The following table shows the colors that are generated for the subnodes by their different alerts:

*Table 72. Alert colors*

| <b>Alert</b>            | <b>Color</b> |
|-------------------------|--------------|
| Probe failed            | Red          |
| RAM increase            | Green        |
| RAM decrease            | Red          |
| Memory increase         | Green        |
| Memory decrease         | Red          |
| New disk                | Green        |
| Missing disk            | Red          |
| Disk defect             | Red          |
| Disk failure predicted  | Red          |
| New filesystem          | Green        |
| Filesystem reconfigured | Red          |
| Scan failed             | Red          |
| Missing filesystem      | Red          |
| Filesystem freespace    | Red          |
| Filesystem free Inodes  | Red          |
| Computer unreachable    | Red          |
| New directory           | Red          |
| Missing directory       | Red          |

Table 72. Alert colors (continued)

| Alert                         | Color |
|-------------------------------|-------|
| Directory quota violation     | Red   |
| Filesystem quota violation    | Red   |
| Computer quota violation      | Red   |
| Network quota violation       | Red   |
| Constraint violation          | Red   |
| Failed job                    | Red   |
| Disk defect quota             | Red   |
| Ping failed                   | Red   |
| Discovery failed              | Red   |
| Aggregator failed             | Red   |
| Filesystem quota failed       | Red   |
| Computer quota failed         | Red   |
| Network quota failed          | Red   |
| Computer discovered           | Red   |
| Report failed                 | Red   |
| DB probe failed               | Red   |
| DB tablespace quota violation | Red   |
| DB instance quota violation   | Red   |
| DB net quota violation        | Red   |
| DB tablespace quota failed    | Red   |
| DB instance quota failed      | Red   |
| DB net quota failed           | Red   |
| Scan failed                   | Red   |
| Archive log overflow          | Red   |
| New tablespace                | Red   |
| Tablespace free space         | Red   |
| Tablespace offline            | Red   |
| Tablespace fragmented         | Red   |
| Tablespace extent size        | Red   |
| Dropped tablespace            | Red   |
| Table fragmented              | Red   |
| Table maximum extents         | Red   |
| Table size                    | Red   |
| Table chained rows            | Red   |
| Table wasted space            | Red   |
| Dropped table                 | Red   |
| NAS discovery failed          | Red   |
| NAS computer discovered       | Green |
| Script failed                 | Red   |
| NetApp quota job failed       | Red   |

Table 72. Alert colors (continued)

| Alert                                  | Color |
|----------------------------------------|-------|
| Device freespace less                  | Red   |
| Device freespace greater               | Red   |
| Dropped device                         | Red   |
| New device                             | Red   |
| DB log freespace                       | Red   |
| DB dump time                           | Red   |
| Disk array discovered                  | Green |
| Disk array missing                     | Red   |
| Filer missing                          | Red   |
| Cache increase                         | Green |
| Cache decrease                         | Red   |
| Automatic filesystem extension         | Red   |
| Stopped automatic filesystem extension | Red   |
| Extend job failed                      | Red   |
| Upgrade failed                         | Red   |
| Added cluster resource group           | Green |
| Removed cluster resource group         | Green |
| Moved cluster resource group           | Green |
| DB report failed                       | Red   |
| File backup failed                     | Red   |
| Hypervisor discovered                  | Green |
| Hypervisor missing                     | Red   |
| Virtual machine discovered             | Green |
| Virtual machine missing                | Red   |

---

## Performance data

Use the IBM Tivoli Storage Productivity Center to view the performance data of devices in your environment.

### General performance guidelines

This topic provides some general performance considerations for configuring the IBM Tivoli Storage Productivity Center environment.

- Do not overpopulate the SLP discovery panel with SLP agent hosts. Remember that Tivoli Storage Productivity Center includes a built-in SLP User Agent (UA) that receives information about SLP Service Agents and Directory Agents (DA) that reside in the same subnet as the Tivoli Storage Productivity Center server.
- Ensure that you have not more than one DA per subnet.
- Misconfiguring the Tivoli Storage Productivity Center Common Information Model Object Manager (CIMOM) discovery preferences might affect performance on auto discovery or on device presence checking. It might also result in application timeouts, as attempts are made to resolve and communicate

with hosts that are unavailable. CIMOMs are not required to communicate with storage systems that use native interfaces. Tivoli Storage Productivity Center uses native interfaces to communicate with IBM System Storage DS8000 Storage Manager, IBM XIV Storage System, and IBM System Storage SAN Volume Controller.

- It should be considered mandatory to run the CIM Agent software on a separate host from the Tivoli Storage Productivity Center server. Attempting to run a full Tivoli Storage Productivity Center implementation on the same host as the CIM agent, can result in dramatically increased wait times for data retrieval. You might also experience resource contention and port conflicts.

## Performance thresholds

Performance thresholds are triggering conditions which are used to monitor a component with user-defined values.

You can monitor the performance of your enterprise by creating alerts on performance thresholds for switches and storage subsystems. By creating alerts that are triggered by performance thresholds, you can be informed about performance issues in your enterprise. Threshold events tell you when a component has fallen outside of the user-defined values. For example, when a threshold value has reached critical stress.

### Threshold boundaries

You can establish your boundaries for the normal expected subsystem performance when defining storage subsystem alerts for performance threshold events. When the collected performance data samples fall outside out of the range you have set, you are notified of this threshold violation so you are aware of the potential problem. The upper boundaries are **Critical Stress** and **Warning Stress**. The lower boundaries are **Warning Idle** and **Critical Idle**. Usually you will want the stress boundaries to be high numbers and the idle to be low numbers. The exception to this rule is **Cache Holding Time Threshold**, where you want the stress numbers to be low and the idle numbers to be high.

If you do not want to be notified of threshold violations for any boundaries, you can leave the boundary field blank and the performance data will not be checked against any value. For example, if the **Critical Idle** and **Warning Idle** fields are left blank, no alerts will be sent for any idle conditions.

The **Ignore triggering condition when the sequential I/O percentage exceeds** check box is active only for the triggering condition **Disk Utilization Percentage Threshold**. It is a filter condition. The default is 80%.

The **Ignore triggering condition when the Back-end Read I/O Rate is less than** check box only applies to the **Back-end Read Response Time** and **Back-end Read Queue Time** thresholds.

The **Ignore triggering condition when the Back-end Write I/O Rate is less than** check box only applies to the **Back-end Write Response Time** and **Back-end Write Queue Time** thresholds.

The **Ignore triggering condition when the Total Back-end I/O Rate is less than** check box only applies to the **Overall Back-end Response Time** threshold.

The **Ignore triggering condition when the Total I/O Rate is less than** check box only applies to the **Non-preferred Node Usage Percentage** threshold.

The **Ignore triggering condition when the Write-cache Delay I/O Rate is less than** check box only applies to the **Write-cache Delay Percentage** threshold.

## Setting the thresholds

Only a few thresholds have defaults and on the other thresholds you will have to determine the best values for stress, idle, critical, and warning values so you can derive the maximum benefit without generating too many false alerts. Since suitable stress thresholds are highly dependant on the type of workload you are running, your exact hardware configuration, the number of physical disks, exact model numbers, and so forth, there are no easy or standard default rules.

One of the best approaches is to monitor your performance for a number of weeks and, using this historical data, determine reasonable values for each threshold setting. Once that is done you can fine tune these settings to minimize the number of false alerts.

### Related tasks:

“Archive and delete files” on page 144

This scenario shows you how to reclaim storage space by archiving and deleting files that meet criteria you define.

### Related reference:

“Controller thresholds” on page 1181

The following table lists and describes the Controller thresholds.

“Port thresholds” on page 1182

Port thresholds are used to set limits for such things as bandwidth utilization, data rates, and I/O operations.

“Array thresholds”

The following table lists and describes the Array thresholds.

## Array thresholds

The following table lists and describes the Array thresholds.

*Table 73. Array thresholds*

| Threshold (Metric)          | Device/Component Type                                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------|--------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Array Thresholds</b>     |                                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Disk Utilization Percentage | TotalStorage Enterprise Storage Server, DS6000, DS8000 array | Sets thresholds on the approximate utilization percentage of the arrays in a particular subsystem; for example, the average percentage of time that the disks associated with the array were busy. The Disk Utilization metric for each array is checked against the threshold boundaries for each collection interval. This threshold is enabled by default for TotalStorage Enterprise Storage Server systems and disabled by default for others. The default threshold boundaries are 80%, 50%, -1, -1. For DS6000 and DS8000 subsystems, this threshold applies only to those ranks which are the only ranks in their associated extent pool. |

Table 73. Array thresholds (continued)

| Threshold (Metric)             | Device/Component Type                                                                                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------------|-----------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Total Back-end I/O Rate        | SAN Volume Controller MDisk group and MDisk, TotalStorage Enterprise Storage Server, DS6000, DS8000 array | Sets thresholds on the average number of I/O operations per second for array and MDisk read and write operations. The Total I/O Rate metric for each array or MDisk is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.                                                                                                                                                                                                                                                                                   |
| Total Back-end Data Rate       | SAN Volume Controller MDisk group and MDisk, TotalStorage Enterprise Storage Server, DS6000, DS8000 array | Sets thresholds on the average number of MB per second that were transferred for array and MDisk read and write operations. The Total Data Rate metric for each array or MDisk is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.                                                                                                                                                                                                                                                                        |
| Back-end Read Response Time    | SAN Volume Controller MDisk, TotalStorage Enterprise Storage Server, DS6000, DS8000 array                 | Sets thresholds on the average number of milliseconds that it took to service each array and MDisk read operation. The Back-end Read Response Time metric for each array or MDisk is checked against the threshold boundaries for each collection interval. Though this threshold is disabled by default, suggested boundary values of 35,25,-1,-1 are pre-populated. A filter is available for this threshold which will ignore any boundary violations if the Back-end Read I/O Rate is less than a specified filter value. The pre-populated filter value is 5.     |
| Back-end Write Response Time   | SAN Volume Controller MDisk, TotalStorage Enterprise Storage Server, DS6000, DS8000 array                 | Sets thresholds on the average number of milliseconds that it took to service each array and MDisk write operation. The Back-end Write Response Time metric for each array or MDisk is checked against the threshold boundaries for each collection interval. Though this threshold is disabled by default, suggested boundary values of 120,80,-1,-1 are pre-populated. A filter is available for this threshold which will ignore any boundary violations if the Back-end Write I/O Rate is less than a specified filter value. The pre-populated filter value is 5. |
| Overall Back-end Response Time | SAN Volume Controller MDisk                                                                               | Sets thresholds on the average number of milliseconds that it took to service each MDisk I/O operation, measured at the MDisk level. The Total Response Time (external) metric for each MDisk is checked against the threshold boundaries for each collection interval. This threshold is disabled by default. A filter is available for this threshold which will ignore any boundary violations if the Total Back-end I/O Rate is less than a specified filter value. The pre-populated filter value is 10.                                                          |



Table 73. Array thresholds (continued)

| Threshold (Metric)        | Device/Component Type       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Back-end Read Queue Time  | SAN Volume Controller MDisk | Sets thresholds on the average number of milliseconds that each read operation spent on the queue before being issued to the back-end device. The Back-end Read Queue Time metric for each MDisk is checked against the threshold boundaries for each collection interval. Though this threshold is disabled by default, suggested boundary values of 5,3,-1,-1 are pre-populated. A filter is available for this threshold which will ignore any boundary violations if the Back-end Read I/O Rate is less than a specified filter value. The pre-populated filter value is 5. Violation of these threshold boundaries means that the SVC deems the MDisk to be overloaded. There is a queue algorithm that determines the number of concurrent I/O operations that the SVC will send to a given MDisk. If there is any queuing (other than during a backup process) then this suggests performance can be improved by resolving the queuing issue.   |
| Back-end Write Queue Time | SAN Volume Controller MDisk | Sets thresholds on the average number of milliseconds that each write operation spent on the queue before being issued to the back-end device. The Back-end Write Queue Time metric for each MDisk is checked against the threshold boundaries for each collection interval. Though this threshold is disabled by default, suggested boundary values of 5,3,-1,-1 are pre-populated. A filter is available for this threshold which will ignore any boundary violations if the Back-end Read I/O Rate is less than a specified filter value. The pre-populated filter value is 5. Violation of these threshold boundaries means that the SVC deems the MDisk to be overloaded. There is a queue algorithm that determines the number of concurrent I/O operations that the SVC will send to a given MDisk. If there is any queuing (other than during a backup process) then this suggests performance can be improved by resolving the queuing issue. |

Table 73. Array thresholds (continued)

| Threshold (Metric)                       | Device/Component Type      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------------|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Peak Back-end Write Response Time        | SAN Volume Controller Node | Sets thresholds on the peak (worst) response time among all MDisk write operations by a node. The Back-end Peak Write Response Time metric for each node is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundary values of 30000,10000,-1,-1. Violation of these threshold boundaries means that the SVC cache is having to "partition-limit" for a given MDisk group. The de-staged data from the SVC cache for this MDisk group is causing the cache to fill up (writes are being received faster than they can be de-staged to disk). If delays reach 30 seconds or more, then the SVC will switch into "short-term mode" where writes are no longer cached for the MDisk Group. |
| Port to Local Node Send Response Time    | SAN Volume Controller Node | Sets thresholds on the average number of milliseconds it took to service each send operation to another node in the local SVC cluster. The Port to Local Node Send Response Time metric for each node is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundary values of 3,1.5,-1,-1. Violation of these threshold boundaries means that it is taking too long to send data between nodes (on the fabric), and suggests that there is either congestion around these FC ports, or an internal SVC microcode problem.                                                                                                                                                                 |
| Port to Local Node Receive Response Time | SAN Volume Controller Node | Sets thresholds on the average number of milliseconds it took to service each receive operation from another node in the local SVC cluster. The Port to Local Node Receive Response Time metric for each node is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundary values of 1,0.5,-1,-1. Violation of these threshold boundaries means that it is taking too long to send data between nodes (on the fabric), and suggests that there is either congestion around these FC ports, or an internal SVC microcode problem.                                                                                                                                                         |

Table 73. Array thresholds (continued)

| Threshold (Metric)                    | Device/Component Type      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------------|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port to Local Node Send Queue Time    | SAN Volume Controller Node | Sets thresholds on the average number of milliseconds that each send operation issued to another node in the local SVC cluster spent on the queue before being issued. The Port to Local Node Send Queue Time metric for each node is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundary values of 2,1,-1,-1. Violation of these threshold boundaries means that the node has to wait too long to send data to other nodes (on the fabric), and suggests congestion on the fabric.              |
| Port to Local Node Receive Queue Time | SAN Volume Controller Node | Sets thresholds on the average number of milliseconds that each receive operation issued to another node in the local SVC cluster spent on the queue before being issued. The Port to Local Node Receive Queue Time metric for each node is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundary values of 1,0.5,-1,-1. Violation of these threshold boundaries means that the node has to wait too long to receive data from other nodes (on the fabric), and suggests congestion on the fabric. |

### Controller thresholds

The following table lists and describes the Controller thresholds.

Table 74. Controller thresholds

| Threshold (Metric)           | Device/Component Type                                                                                    | Description                                                                                                                                                                                                                                                                                                                             |
|------------------------------|----------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Controller Thresholds</b> |                                                                                                          |                                                                                                                                                                                                                                                                                                                                         |
| Total I/O Rate (overall)     | TotalStorage Enterprise Storage Server, DS6000, DS8000 controller<br><br>SAN Volume Controller I/O group | Sets threshold on the average number of I/O operations per second for read and write operations, for the subsystem controllers (clusters) or I/O groups. The Total I/O Rate metric for each controller or I/O group is checked against the threshold boundaries for each collection interval. These thresholds are disabled by default. |
| Total Data Rate              | TotalStorage Enterprise Storage Server, DS6000, DS8000 controller<br><br>SAN Volume Controller I/O group | Sets threshold on the average number of MB per second for read and write operations for the subsystem controllers (clusters) or I/O groups. The Total Data Rate metric for each controller or I/O group is checked against the threshold boundaries for each collection interval. These thresholds are disabled by default.             |

Table 74. Controller thresholds (continued)

| Threshold (Metric)                  | Device/Component Type                                                                         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------------------------------|-----------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NVS Full Percentage                 | TotalStorage Enterprise Storage Server, DS6000, DS8000 controller                             | Sets thresholds on the percentage of time that NVS space constraints caused I/O operations to be delayed, for the subsystem controllers (clusters). The NVS Full Percentage metric for each controller is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundaries of 10, 3, -1, -1.                                                                                                                                                                             |
| Cache Holding Time                  | TotalStorage Enterprise Storage Server, DS6000, DS8000 controller                             | Sets thresholds on the average cache holding time, in seconds, for I/O data in the subsystem controllers (clusters). Shorter time periods indicate adverse performance. The Cache Holding Time metric for each controller is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundaries of 30, 60, -1, -1.                                                                                                                                                         |
| Write-cache Delay Percentage        | TotalStorage Enterprise Storage Server, DS6000, DS8000 controller, SAN Volume Controller node | Sets thresholds on the percentage of I/O operations that were delayed due to write-cache space constraints. This metric for each controller or node is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundaries of 10, 3, -1, -1. In addition, a filter is available for this threshold which will ignore any boundary violations if the Write-cache Delay I/O Rate is less than a specified filter value. The pre-populated filter value is 10 I/Os per second. |
| Non-Preferred Node Usage Percentage | SAN Volume Controller I/O group                                                               | Sets thresholds on the Non-Preferred Node Usage Percentage of an I/O group. This metric of each I/O group is checked against the threshold boundaries at each collection interval. This threshold is disabled by default. In addition, a filter is available for this threshold which will ignore any boundary violations if the Total I/O Rate of the I/O group is less than a specified filter value.                                                                                                                                      |

## Port thresholds

Port thresholds are used to set limits for such things as bandwidth utilization, data rates, and I/O operations.

Table 75. Port thresholds

| Threshold (Metric)     | Device/Component Type | Description |
|------------------------|-----------------------|-------------|
| <b>Port Thresholds</b> |                       |             |

Table 75. Port thresholds (continued)

| Threshold (Metric)         | Device/Component Type                                                   | Description                                                                                                                                                                                                                                                                                 |
|----------------------------|-------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Total Port I/O Rate        | TotalStorage Enterprise Storage Server, DS6000, DS8000 switch port, XIV | Sets thresholds for ports on the average number of I/O operations or packets per second for send and receive operations. The Total I/O Rate metric for each port is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.           |
| Total Port Data Rate       | TotalStorage Enterprise Storage Server, DS6000, DS8000 switch port, XIV | Sets thresholds for ports on the average number of MB per second for send and receive operations. The Total Data Rate metric for each port is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.                                 |
| Overall Port Response Time | TotalStorage Enterprise Storage Server, DS6000, DS8000                  | Sets thresholds for ports on the average number of milliseconds that it takes to service each send and receive I/O operation. The Total Response Time metric for each port is checked against the threshold boundaries for each collection interval. This threshold is disabled by default. |
| Error Frame Rate           | DS8000                                                                  | Sets thresholds on the average number of frames per second received in error by ports. The Error Frame Rate metric for each port is checked against the threshold boundary for each collection interval. This threshold is disabled by default.                                             |
| Link Failure Rate          | DS8000 SAN Volume Controller                                            | Sets thresholds on the average number of link errors per second for ports. The Link Failure Rate metric for each port is checked against the threshold boundary for each collection interval. The default boundaries are 0.0030,-1,-1,-1 cnt/s. This threshold is disabled by default.      |

Table 75. Port thresholds (continued)

| Threshold (Metric)                  | Device/Component Type                                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------------------------|--------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Total Port Data Rate                | Switch port                                            | Sets thresholds on the critical and warning data rates for stress and idle in MB per second. The Total Port Data Rate metric for each port is checked against the threshold boundary for each collection interval. This threshold is disabled by default.                                                                                                                                                                                                                                                                            |
| Total Port Packet Rate              | Switch port                                            | Sets thresholds on critical and warning data rates for stress and idle conditions in packets per second. For example, a critical stress or warning stress condition occurs when the upper boundary for the packet rate of a switch is detected. A critical idle or warning idle condition occurs when the lower boundary for the packet rate of a switch is detected. The Total Port Packet Rate metric for each port is checked against the threshold boundary for each collection interval. This threshold is disabled by default. |
| Port Send Utilization Percentage    | TotalStorage Enterprise Storage Server, DS6000, DS8000 | Sets thresholds on the average amount of time that ports are busy sending data. The metric for each port is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.                                                                                                                                                                                                                                                                                                            |
| Port Receive Utilization Percentage | TotalStorage Enterprise Storage Server, DS6000, DS8000 | Sets thresholds on the average amount of time that ports are busy receiving data. The metric for each port is checked against the threshold boundaries for each collection interval. This threshold is disabled by default.                                                                                                                                                                                                                                                                                                          |

Table 75. Port thresholds (continued)

| Threshold (Metric)                | Device/Component Type                                                                  | Description                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------|----------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port Send Bandwidth Percentage    | TotalStorage Enterprise Storage Server, DS8000, SAN Volume Controller switch port, XIV | Sets thresholds on the average port bandwidth utilization percentage for send operations. The Port Send Utilization Percentage metric is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundaries 85,75,-1,-1.                                                                                   |
| Port Receive Bandwidth Percentage | TotalStorage Enterprise Storage Server, DS8000, SAN Volume Controller switch port, XIV | Sets thresholds on the average port bandwidth utilization percentage for receive operations. The Port Send Utilization Percentage metric is checked against the threshold boundaries for each collection interval. This threshold is enabled by default, with default boundaries 85,75,-1,-1.                                                                                |
| CRC Error Rate                    | DS8000 port, SAN Volume Controller switch port                                         | Sets thresholds on the average number of frames received in which the cyclic redundancy check (CRC) in a frame does not match the CRC computed by the receiver. The CRC Error Rate metric for each port is checked against the threshold boundary for each collection interval. The default boundaries are 0.033, 0.01, -1, -1 cnt/s. This threshold is disabled by default. |
| Invalid Transmission Word Rate    | DS8000, SAN Volume Controller switch port                                              | Sets thresholds on the average number of bit errors detected on a port. The Invalid Transmission Word Rate metric for each port is checked against the threshold boundary for each collection interval. The default boundaries are 0.033, 0.01, -1, -1 cnt/s. This threshold is disabled by default.                                                                         |

Table 75. Port thresholds (continued)

| Threshold (Metric)                | Device/Component Type | Description                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Zero Buffer - Buffer Credit Timer | SAN Volume Controller | Sets thresholds on the number of microseconds that a port has been unable to send frames because of a lack of buffer credit since the last node reset. The Zero Buffer-Buffer Credit Timer metric for each port is checked against the threshold boundary for each collection interval. The default boundaries are 198000000,75000000,-1,-1 us. This threshold is disabled by default. |

## Performance metrics

A description of the metrics that are displayed in the columns of performance reports is provided.

### Common columns

The following table contains information about the columns that are common among performance reports.

| Column   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Time     | Date and time that the data was collected                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Interval | <p>Size of the sample interval in seconds. You can specify a minimum interval length of five minutes and a maximum interval length of sixty minutes for the following models:</p> <ul style="list-style-type: none"> <li>• TotalStorage Enterprise Storage Server</li> <li>• DS6000</li> <li>• DS8000</li> <li>• the XIV system</li> </ul> <p>For SAN Volume Controller models earlier than V4.1, you can specify a minimum interval length of 15 minutes and a maximum interval length of sixty minutes. For SAN Volume Controller models V4.1 and later, you can specify a minimum interval length of 5 minutes, and a maximum interval length of sixty minutes.</p> |

**Note:** When you view metrics for the ESS, DS6000, and DS8000 storage systems, you must take into account the following differences between IBM Tivoli Storage Productivity Center reports and the native reports of those systems:

- Tivoli Storage Productivity Center reports display port performance metrics as send and receive metrics (for example, **Send Data Rate** and **Receive Data Rate**).
- Storage system native reports (for example, reports based on data collected by the DS CLI) display port performance metrics as read and write metrics (for example, **Bytread** and **Bytewrite**).

When a host performs a read operation, the DS6000 and DS8000 ports send data to the host. Therefore "read" metrics in DS6000 and DS8000 reports correspond to



"send" metrics in Tivoli Storage Productivity Center reports. When a host performs a write operation, DS6000. and DS8000 ports receive data from the host. Therefore "write" metrics in DS6000. and DS8000 reports correspond to "receive" metrics in Tivoli Storage Productivity Center reports.

When you view port Peer-to-Peer Remote Copy (PPRC) performance metrics, you must take into account the following additional differences between Tivoli Storage Productivity Center reports and native reports for storage systems:

- Metrics for PPRC reads in storage system native reports are represented as PPRC receives in Tivoli Storage Productivity Center (reads = receives).
- Metrics for PPRC writes in storage system native reports are represented as PPRC sends in Tivoli Storage Productivity Center (writes = sends).

## XIV system metrics

To distinguish between metrics that were introduced in different versions of the XIV systems, the following convention was used in this topic:

- <sup>1</sup> is displayed next to metrics that are available in the XIV system version 10.2.2 or later.
- <sup>2</sup> is displayed next to metrics that are available in the XIV system version 10.2.4 or later.

For example:

- The Read I/O Rate (overall) metric is available for the XIV systems version 10.2.2 and later. In the **Devices: components** column of the list of metrics, the entry for Read I/O Rate (overall) is displayed like this: XIV<sup>1</sup>: volume, module, subsystem
- The Small Transfers Response Time metric is available for the XIV systems version 10.2.4 and later. In the **Devices: components** column of the list of metrics, the entry for Small Transfers Response Time is displayed like this: XIV<sup>2</sup>: volume, module, subsystem

## Volume-based metrics

The following table contains information about volume-based metrics.

**Note:** Tivoli Storage Productivity Center does not calculate volume-based metrics if there are space efficient volumes allocated in an extent pool consisting of multiple ranks. In this case, the columns for volume-based metrics display the value **N/A** in the Storage Subsystem Performance > By Array report for the arrays associated with that extent pool. However, if there are no space efficient volumes allocated in a multi-rank extent pool, or if the space efficient volumes are allocated in an extent pool consisting of a single rank, then this limitation does not apply and all volume-based metrics are displayed in the By Array report.

| Column                     | Devices: components                                     | Description                                                                                                                   |
|----------------------------|---------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>I/O Rates</b>           |                                                         |                                                                                                                               |
| Read I/O Rate (normal)     | ESS/DS6000/DS8000: volume, array, controller, subsystem | Average number of I/O operations per second for nonsequential read operations for a component over a specified time interval. |
| Read I/O rate (sequential) | ESS/DS6000/DS8000: volume, array, controller, subsystem | Average number of I/O operations per second for sequential read operations for a component over a specified time interval.    |

| Column                       | Devices: components                                                                                                                                                                                                    | Description                                                                                                                                                 |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Read I/O Rate (overall)      | ESS/DS6000/DS8000: volume, array, controller, subsystem<br><br>SVC: VDisk, node, I/O group, MDisk group, subsystem<br><br>SMI-S BSP: volume, controller, subsystem<br><br>XIV <sup>1</sup> : volume, module, subsystem | Average number of I/O operations per second for both sequential and nonsequential read operations for a component over a specified time interval.           |
| Write I/O Rate (normal)      | ESS/DS6000/DS8000: volume, array, controller, subsystem                                                                                                                                                                | Average number of I/O operations per second for nonsequential write operations for a component over a specified time interval.                              |
| Write I/O Rate (sequential)  | ESS/DS6000/DS8000: volume, array, controller, subsystem                                                                                                                                                                | Average number of I/O operations per second for sequential write operations for a component over a specified time interval.                                 |
| Write I/O Rate (overall)     | ESS/DS6000/DS8000: volume, array, controller, subsystem<br><br>SVC: VDisk, node, I/O group, MDisk group, subsystem<br><br>SMI-S BSP: volume, controller, subsystem<br><br>XIV <sup>1</sup> : volume, module, subsystem | Average number of I/O operations per second for both sequential and nonsequential write operations for a component over a specified time interval.          |
| Total I/O Rate (normal)      | ESS/DS6000/DS8000: volume, array, controller, subsystem                                                                                                                                                                | Average number of I/O operations per second for nonsequential read and write operations for a component over a specified time interval.                     |
| Total I/O Rate (sequential)  | ESS/DS6000/DS8000: volume, array, controller, subsystem                                                                                                                                                                | Average number of I/O operations per second for sequential read and write operations for a component over a specified time interval.                        |
| Total I/O Rate (overall)     | ESS/DS6000/DS8000: volume, array, controller, subsystem<br><br>SVC: VDisk, node, I/O group, MDisk group, subsystem<br><br>SMI-S BSP: volume, controller, subsystem<br><br>XIV <sup>1</sup> : volume, module, subsystem | Average number of I/O operations per second for both sequential and nonsequential read and write operations for a component over a specified time interval. |
| Global Mirror Write I/O Rate | SVC: VDisk, node, I/O group, subsystem                                                                                                                                                                                 | Average number of write operations per second issued to the Global Mirror secondary site for a component over a specified time interval.                    |

| Column                                     | Devices: components                                     | Description                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------------------------|---------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Global Mirror Overlapping Write Percentage | SVC: VDisk, node, I/O group, subsystem                  | Average percentage of write operations issued by the Global Mirror primary site which were serialized overlapping writes for a component over a specified time interval. For SVC 4.3.1 and later, some overlapping writes are processed in parallel (are not serialized) and are excluded. For earlier SVC versions, all overlapping writes were serialized. |
| Global Mirror Overlapping Write I/O Rate   | SVC: VDisk, node, I/O group, subsystem                  | Average number of serialized overlapping write operations per second encountered by the Global Mirror primary site for a component over a specified time interval. For SVC 4.3.1 and later, some overlapping writes are processed in parallel (are not serialized) and are excluded. For earlier SVC versions, all overlapping writes are serialized.        |
| HPF Read I/O Rate                          | DS8000: volume, array, controller, subsystem            | Average number of read operations per second that were issued via the High Performance FICON® (HPF) feature of the storage subsystem for a component over a specified time interval.                                                                                                                                                                         |
| HPF Write I/O Rate                         | DS8000: volume, array, controller, subsystem            | Average number of write operations per second that were issued via the High Performance FICON (HPF) feature of the storage subsystem for a component over a specified time interval.                                                                                                                                                                         |
| Total HPF I/O Rate                         | DS8000: volume, array, controller, subsystem            | Average number of read and write operations per second that were issued via the High Performance FICON (HPF) feature of the storage subsystem for a component over a specified time interval.                                                                                                                                                                |
| HPF I/O Percentage                         | DS8000: volume, array, controller, subsystem            | The percentage of all I/O operations that were issued via the High Performance FICON (HPF) feature of the storage subsystem for a component over a specified time interval.                                                                                                                                                                                  |
| PPRC Transfer Rate                         | ESS/DS6000/DS8000: volume, array, controller, subsystem | Average number of track transfer operations per second for PPRC usage for a component over a specified time interval. This metric shows the activity for the source of the PPRC relationship, but shows no activity for the target.                                                                                                                          |
| Small Transfers I/O Percentage             | XIV <sup>2</sup> : volume, module, subsystem            | Percentage of I/O operations over a specified interval. Applies to data transfer sizes that are <= 8 KB.                                                                                                                                                                                                                                                     |
| Medium Transfers I/O Percentage            | XIV <sup>2</sup> : volume, module, subsystem            | Percentage of I/O operations over a specified interval. Applies to data transfer sizes that are > 8 KB and <= 64 KB.                                                                                                                                                                                                                                         |
| Large Transfers I/O Percentage             | XIV <sup>2</sup> : volume, module, subsystem            | Percentage of I/O operations over a specified interval. Applies to data transfer sizes that are > 64 KB and <= 512 KB.                                                                                                                                                                                                                                       |
| Very Large Transfers I/O Percentage        | XIV <sup>2</sup> : volume, module, subsystem            | Percentage of I/O operations over a specified interval. Applies to data transfer sizes that are > 512 KB.                                                                                                                                                                                                                                                    |
| <b>Cache hit percentages</b>               |                                                         |                                                                                                                                                                                                                                                                                                                                                              |

| Column                                   | Devices: components                                                                                                                                                                                       | Description                                                                                                                     |
|------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| Read Cache Hits Percentage (normal)      | ESS/DS6000/DS8000: volume, array, controller, subsystem                                                                                                                                                   | Percentage of cache hits for nonsequential read operations for a component over a specified time interval.                      |
| Read Cache Hits Percentage (sequential)  | ESS/DS6000/DS8000: volume, array, controller, subsystem                                                                                                                                                   | Percentage of cache hits for sequential read operations for a component over a specified time interval.                         |
| Read Cache Hits Percentage (overall)     | ESS/DS6000/DS8000: volume, array, controller, subsystem<br><br>SVC: VDisk, node, I/O group, subsystem<br><br>SMI-S BSP: volume, controller, subsystem<br><br>XIV <sup>1</sup> : volume, module, subsystem | Percentage of cache hits for both sequential and nonsequential read operations for a component over a specified time interval.  |
| Write Cache Hits Percentage (normal)     | ESS/DS6000/DS8000: volume, array, controller, subsystem                                                                                                                                                   | Percentage of cache hits for nonsequential write operations for a component over a specified time interval.                     |
| Write Cache Hits Percentage (sequential) | ESS/DS6000/DS8000: volume, array, controller, subsystem                                                                                                                                                   | Percentage of cache hits for sequential write operations for a component over a specified time interval.                        |
| Write Cache Hits Percentage (overall)    | ESS/DS6000/DS8000: volume, array, controller, subsystem<br><br>SVC: VDisk, node, I/O group, subsystem<br><br>SMI-S BSP: volume, controller, subsystem<br><br>XIV <sup>1</sup> : volume, module, subsystem | Percentage of cache hits for both sequential and nonsequential write operations for a component over a specified time interval. |
| Total Cache Hits Percentage (normal)     | ESS/DS6000/DS8000: volume, array, controller, subsystem                                                                                                                                                   | Percentage of cache hits for nonsequential read and write operations for a component over a specified time interval.            |
| Total Cache Hits Percentage (sequential) | ESS/DS6000/DS8000: volume, array, controller, subsystem                                                                                                                                                   | Percentage of cache hits for sequential read and write operations for a component over a specified time interval.               |

| Column                                | Devices: components                                                                                                                                                                                                    | Description                                                                                                                                               |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Total Cache Hits Percentage (overall) | ESS/DS6000/DS8000: volume, array, controller, subsystem<br><br>SVC: VDisk, node, I/O group, subsystem<br><br>SMI-S BSP: volume, controller, subsystem<br><br>XIV <sup>1</sup> : volume, module, subsystem              | Percentage of cache hits for both sequential and nonsequential read and write operations for a component over a specified time interval.                  |
| Readahead Percentage of Cache Hits    | SVC: VDisk, node, I/O group, subsystem                                                                                                                                                                                 | Percentage of all read cache hits which occurred on prestaged data.                                                                                       |
| Dirty Write Percentage of Cache Hits  | SVC: VDisk, node, I/O group, subsystem                                                                                                                                                                                 | Percentage of all write cache hits which occurred on already dirty data in the cache.                                                                     |
| Read Data Cache Hit Percentage        | XIV <sup>2</sup> : volume, module, subsystem                                                                                                                                                                           | Percentage of read data that was read from the cache over a specified time interval.                                                                      |
| Write Data Cache Hit Percentage       | XIV <sup>2</sup> : volume, module, subsystem                                                                                                                                                                           | Percentage of write data that was written to the cache over a specified time interval.                                                                    |
| Total Data Cache Hit Percentage       | XIV <sup>2</sup> : volume, module, subsystem                                                                                                                                                                           | Percentage of all data that was read from or written to the cache for a component over a specified time interval.                                         |
| <b>Data rates</b>                     |                                                                                                                                                                                                                        |                                                                                                                                                           |
| Read Data Rate                        | ESS/DS6000/DS8000 volume, array, controller, subsystem<br><br>SVC: VDisk, node, I/O group, MDisk group, subsystem<br><br>SMI-S BSP: volume, controller, subsystem<br><br>XIV <sup>1</sup> : volume, module, subsystem  | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for read operations for a component over a specified time interval.  |
| Write Data Rate                       | ESS/DS6000/DS8000: volume, array, controller, subsystem<br><br>SVC: VDisk, node, I/O group, MDisk group, subsystem<br><br>SMI-S BSP: volume, controller, subsystem<br><br>XIV <sup>1</sup> : volume, module, subsystem | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for write operations for a component over a specified time interval. |

| Column                           | Devices: components                                                                                                                                                                                                    | Description                                                                                                                                                        |
|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Total Data Rate                  | ESS/DS6000/DS8000: volume, array, controller, subsystem<br><br>SVC: VDisk, node, I/O group, MDisk group, subsystem<br><br>SMI-S BSP: volume, controller, subsystem<br><br>XIV <sup>1</sup> : volume, module, subsystem | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for read and write operations for a component over a specified time interval. |
| Small Transfers Data Percentage  | XIV <sup>2</sup> : volume, module, subsystem                                                                                                                                                                           | Percentage of data that was transferred over a specified interval. Applies to I/O operations with data transfer sizes that are <= 8 KB.                            |
| Medium Transfers Data Percentage | XIV <sup>2</sup> : volume, module, subsystem                                                                                                                                                                           | Percentage of data that was transferred over a specified interval. Applies to I/O operations with data transfer sizes that are > 8 KB and <= 64 KB.                |
| Large Transfers Data Percentage  | XIV <sup>2</sup> : volume, module, subsystem                                                                                                                                                                           | Percentage of data that was transferred over a specified interval. Applies to I/O operations with data transfer sizes that are > 64 KB and <= 512 KB.              |
| Very Large Transfers Data        | XIV <sup>2</sup> : volume, module, subsystem                                                                                                                                                                           | Percentage of data that was transferred over a specified interval. Applies to I/O operations with data transfer sizes that are > 512 KB.                           |
| <b>Response times</b>            |                                                                                                                                                                                                                        |                                                                                                                                                                    |
| Read Response Time               | ESS/DS6000/DS8000: volume, array, controller, subsystem<br><br>SVC: VDisk, node, I/O group, MDisk group, subsystem<br><br>SMI-S BSP: volume, controller, subsystem<br><br>XIV <sup>1</sup> : volume, module, subsystem | Average number of milliseconds that it took to service each read operation for a component over a specified time interval.                                         |
| Write Response Time              | ESS/DS6000/DS8000: volume, array, controller, subsystem<br><br>SVC: VDisk, node, I/O group, MDisk group, subsystem<br><br>SMI-S BSP volume, controller, subsystem<br><br>XIV <sup>1</sup> : volume, module, subsystem  | Average number of milliseconds that it took to service each write operation for a component over a specified time interval.                                        |

| Column                                           | Devices: components                                                                                                                                                                                                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Overall Response Time                            | ESS/DS6000/DS8000: volume, array, controller, subsystem<br><br>SVC: VDisk, node, I/O group, MDisk group, subsystem<br><br>SMI-S BSP: volume, controller, subsystem<br><br>XIV <sup>1</sup> : volume, module, subsystem | Average number of milliseconds that it took to service each I/O operation (read and write) for a component over a specified time interval.                                                                                                                                                                                                                                                                                                                                                  |
| Peak Read Response Time                          | SVC: VDisk, node, I/O group, subsystem                                                                                                                                                                                 | The peak (worst) response time among all read operations.                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Peak Write Response Time                         | SVC: VDisk, node, I/O group, subsystem                                                                                                                                                                                 | The peak (worst) response time among all write operations.                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Global Mirror Write Secondary Lag                | SVC: VDisk, node, I/O group, subsystem                                                                                                                                                                                 | The average number of additional milliseconds it takes to service each secondary write operation for Global Mirror, over and above the time that is required to service primary writes.                                                                                                                                                                                                                                                                                                     |
| Overall Host Attributed Response Time Percentage | SVC: VDisk, node, I/O group, subsystem                                                                                                                                                                                 | The percentage of the average response time, both read response time and write response time, that can be attributed to delays from host systems. This metric is provided to help diagnose slow hosts and poorly performing fabrics. The value is based on the time taken for hosts to respond to transfer-ready notifications from the SVC nodes (for read) and the time taken for hosts to send the write data after the node has responded to a transfer-ready notification (for write). |
| Read Cache Hit Response Time                     | XIV <sup>2</sup> : volume, module, subsystem                                                                                                                                                                           | Average number of milliseconds that it takes to service each read cache hit operation over a specified time interval.                                                                                                                                                                                                                                                                                                                                                                       |
| Write Cache Hit Response Time                    | XIV <sup>2</sup> : volume, module, subsystem                                                                                                                                                                           | Average number of milliseconds that it takes to service each write cache hit operation over a specified time interval.                                                                                                                                                                                                                                                                                                                                                                      |
| Overall Cache Hit Response Time                  | XIV <sup>2</sup> : volume, module, subsystem                                                                                                                                                                           | Average number of milliseconds that it takes to service each read cache hit operation and each write cache hit operation over a specified time interval.                                                                                                                                                                                                                                                                                                                                    |
| Read Cache Miss Response Time                    | XIV <sup>2</sup> : volume, module, subsystem                                                                                                                                                                           | Average number of milliseconds that it takes to service each read cache miss operation over a specified time interval.                                                                                                                                                                                                                                                                                                                                                                      |
| Write Cache Miss Response Time                   | XIV <sup>2</sup> : volume, module, subsystem                                                                                                                                                                           | Average number of milliseconds that it takes to service each write cache miss operation over a specified time interval.                                                                                                                                                                                                                                                                                                                                                                     |

| Column                             | Devices: components                                                                                                                                                                                                    | Description                                                                                                                                                |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Overall Cache Miss Response Time   | XIV <sup>2</sup> : volume, module, subsystem                                                                                                                                                                           | Average number of milliseconds that it takes to service each read cache miss operation and each write cache miss operation over a specified time interval. |
| Small Transfers Response Time      | XIV <sup>2</sup> : volume, module, subsystem                                                                                                                                                                           | Average number of milliseconds that it takes to service each I/O operation. Applies to data transfer sizes that are <= 8 KB.                               |
| Medium Transfers Response Time     | XIV <sup>2</sup> : volume, module, subsystem                                                                                                                                                                           | Average number of milliseconds that it takes to service each I/O operation. Applies to data transfer sizes that are > 8 KB and <= 64 KB.                   |
| Large Transfers Response Time      | XIV <sup>2</sup> : volume, module, subsystem                                                                                                                                                                           | Average number of milliseconds that it takes to service each I/O operation. Applies to data transfer sizes that are > 64 KB and <= 512 KB.                 |
| Very Large Transfers Response Time | XIV <sup>2</sup> : volume, module, subsystem                                                                                                                                                                           | Average number of milliseconds that it takes to service each I/O operation. Applies to data transfer sizes that are > 512 KB.                              |
| <b>Transfer sizes</b>              |                                                                                                                                                                                                                        |                                                                                                                                                            |
| Read Transfer Size                 | ESS/DS6000/DS8000: volume, array, controller, subsystem<br><br>SVC: VDisk, node, I/O group, MDisk group, subsystem<br><br>SMI-S BSP: volume, controller, subsystem<br><br>XIV <sup>1</sup> : volume, module, subsystem | Average number of KB per I/O for read operations.                                                                                                          |
| Write Transfer Size                | ESS/DS6000/DS8000: volume, array, controller, subsystem<br><br>SVC: VDisk, node, I/O group, MDisk group, subsystem<br><br>SMI-S BSP: volume, controller, subsystem<br><br>XIV <sup>1</sup> : volume, module, subsystem | Average number of KB per I/O for write operations.                                                                                                         |



| Column                               | Devices: components                                                                                                                                                                                                    | Description                                                                                                                                                                                              |
|--------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Overall Transfer Size                | ESS/DS6000/DS8000: volume, array, controller, subsystem<br><br>SVC: VDisk, node, I/O group, MDisk group, subsystem<br><br>SMI-S BSP: volume, controller, subsystem<br><br>XIV <sup>1</sup> : volume, module, subsystem | Average number of KB per I/O for read and write operations.                                                                                                                                              |
| <b>Write-cache constraints</b>       |                                                                                                                                                                                                                        |                                                                                                                                                                                                          |
| Write-cache Delay Percentage         | ESS/DS6000/DS8000: volume, array, controller, subsystem<br><br>SVC: VDisk, node, I/O group, subsystem                                                                                                                  | Percentage of I/O operations that were delayed due to write-cache space constraints or other conditions for a component over a specified time interval. (The ratio of delayed operations to total I/Os.) |
| Write-cache Delayed I/O Rate         | ESS/DS6000/DS8000: volume, array, controller, subsystem<br><br>SVC: VDisk, node, I/O group, subsystem                                                                                                                  | Average number of I/O operations per second that were delayed due to write-cache space constraints or other conditions for a component over a specified time interval.                                   |
| Write-cache Overflow Percentage      | SVC: VDisk, node, I/O group, subsystem                                                                                                                                                                                 | Percentage of write operations that were delayed due to lack of write-cache space for a component over a specified time interval.                                                                        |
| Write-cache Overflow I/O Rate        | SVC: VDisk, node, I/O group, subsystem                                                                                                                                                                                 | Average number of tracks per second that were delayed due to lack of write-cache space for a component over a specified time interval.                                                                   |
| Write-cache Flush-through Percentage | SVC: VDisk, node, I/O group, subsystem                                                                                                                                                                                 | Percentage of write operations that were processed in Flush-through write mode for a component over a specified time interval.                                                                           |
| Write-cache Flush-through I/O Rate   | SVC: VDisk, node, I/O group, subsystem                                                                                                                                                                                 | Average number of tracks per second that were processed in Flush-through write mode for a component over a specified time interval.                                                                      |
| Write-cache Write-through Percentage | SVC: VDisk, node, I/O group, subsystem                                                                                                                                                                                 | Percentage of write operations that were processed in Write-through write mode for a component over a specified time interval.                                                                           |
| Write-cache Write-through I/O Rate   | SVC: VDisk, node, I/O group, subsystem                                                                                                                                                                                 | Average number of tracks per second that were processed in Write-through write mode for a component over a specified time interval.                                                                      |
| <b>Record mode reads</b>             |                                                                                                                                                                                                                        |                                                                                                                                                                                                          |
| Record Mode Read I/O Rate            | ESS/DS6000/DS8000: volume, array, controller                                                                                                                                                                           | Average number of I/O operations per second for record mode read operations for a component over a specified time interval.                                                                              |

| Column                               | Devices: components                                                                           | Description                                                                                                                                          |
|--------------------------------------|-----------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| Record Mode Read Cache %             | ESS/DS6000/DS8000:<br>volume, array, controller                                               | Percentage of cache hits for record mode read operations for a component over a specified time interval.                                             |
| <b>Cache transfers</b>               |                                                                                               |                                                                                                                                                      |
| Disk to Cache I/O Rate               | ESS/DS6000/DS8000:<br>volume, array, controller<br><br>SVC: VDisk, node, I/O group, subsystem | Average number of I/O operations (track transfers) per second for disk to cache transfers for a component over a specified time interval.            |
| Cache to Disk I/O Rate               | ESS/DS6000/DS8000:<br>volume, array, controller<br><br>SVC: VDisk, node, I/O group, subsystem | Average number of I/O operations (track transfers) per second for cache to disk transfers for a component over a specified time interval.            |
| <b>Miscellaneous computed values</b> |                                                                                               |                                                                                                                                                      |
| Cache Holding Time                   | ESS/DS6000/DS8000:<br>controller, subsystem                                                   | Average cache holding time, in seconds, for I/O data in this subsystem controller (cluster). Shorter time periods indicate adverse performance.      |
| CPU Utilization                      | SVC: node, I/O group, subsystem                                                               | Average utilization percentage of the processors.                                                                                                    |
| Non-Preferred Node Usage Percentage  | SVC: VDisk, I/O group                                                                         | The overall percentage of I/O performed or data transferred via the non-preferred nodes of the VDIs, for a component over a specified time interval. |
| Volume Utilization                   | ESS/DS6000/DS8000:<br>volume<br><br>SVC: VDisk<br><br>XIV <sup>1</sup> : volume               | The approximate utilization percentage of a volume over a specified time interval (the average percent of time that the volume was busy).            |

## Back-end-based metrics

The following table contains information about back-end-based metrics.

| Column                 | Devices: components                                                                                                 | Description                                                       |
|------------------------|---------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------|
| <b>I/O rates</b>       |                                                                                                                     |                                                                   |
| Backend Read I/O Rate  | ESS/DS6000/DS8000:<br>rank, array, controller, subsystem<br><br>SVC: node, I/O group, MDisk, MDisk group, subsystem | Average number of I/O operations per second for read operations.  |
| Backend Write I/O Rate | ESS/DS6000/DS8000:<br>rank, array, controller, subsystem<br><br>SVC: node, I/O group, MDisk, MDisk group, subsystem | Average number of I/O operations per second for write operations. |

| Column                        | Devices: components                                                                                                          | Description                                                                                                                                                                                        |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Total Backend I/O Rate        | ESS/DS6000/DS8000:<br>rank, array, controller,<br>subsystem<br><br>SVC: node, I/O group,<br>MDisk, MDisk group,<br>subsystem | Average number of I/O operations per second for read and write operations.                                                                                                                         |
| <b>Data rates</b>             |                                                                                                                              |                                                                                                                                                                                                    |
| Backend Read Data Rate        | ESS/DS6000/DS8000:<br>rank, array, controller,<br>subsystem<br><br>SVC: node, I/O group,<br>MDisk, MDisk group,<br>subsystem | Average number of megabytes (2 <sup>20</sup> bytes) that were transferred for read operations.                                                                                                     |
| Backend Write Data Rate       | ESS/DS6000/DS8000:<br>rank, array, controller,<br>subsystem<br><br>SVC: node, I/O group,<br>MDisk, MDisk group,<br>subsystem | Average number of megabytes (2 <sup>20</sup> bytes) that were transferred for write operations.                                                                                                    |
| Total Backend Data Rate       | ESS/DS6000/DS8000:<br>rank, array, controller,<br>subsystem<br><br>SVC: node, I/O group,<br>MDisk, MDisk group,<br>subsystem | Average number of megabytes (2 <sup>20</sup> bytes) that were transferred for read and write operations.                                                                                           |
| <b>Response times</b>         |                                                                                                                              |                                                                                                                                                                                                    |
| Backend Read Response Time    | ESS/DS6000/DS8000:<br>rank, array, controller,<br>subsystem<br><br>SVC: node, I/O group,<br>MDisk, MDisk group,<br>subsystem | Average number of milliseconds that it took to respond to each read operation. For SAN Volume Controller models, this is the external response time of the managed disks (MDisks).                 |
| Backend Write Response Time   | ESS/DS6000/DS8000:<br>rank, array, controller,<br>subsystem<br><br>SVC: node, I/O group,<br>MDisk, MDisk group,<br>subsystem | Average number of milliseconds that it took to respond to each write operation. For SAN Volume Controller models, this is the external response time of the managed disks (MDisks).                |
| Overall Backend Response Time | ESS/DS6000/DS8000:<br>rank, array, controller,<br>subsystem<br><br>SVC: node, I/O group,<br>MDisk, MDisk group,<br>subsystem | Average number of milliseconds that it took to respond to each I/O operation (read and write). For SAN Volume Controller models, this is the external response time of the managed disks (MDisks). |
| Backend Read Queue Time       | SVC: node, I/O group,<br>MDisk, MDisk group,<br>subsystem                                                                    | Average number of milliseconds that each read operation spent on the queue before being issued to the back-end device.                                                                             |

| <b>Column</b>                     | <b>Devices: components</b>                                                                                       | <b>Description</b>                                                                                                                                                                                                                               |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Backend Write Queue Time          | SVC: node, I/O group, MDisk, MDisk group, subsystem                                                              | Average number of milliseconds that each write operation spent on the queue before being issued to the back-end device.                                                                                                                          |
| Overall Backend Queue Time        | SVC: node, I/O group, MDisk, MDisk group, subsystem                                                              | Average number of milliseconds that read and write operations spent on the queue before being issued to the back-end device.                                                                                                                     |
| Peak Back-end Read Response Time  | SVC: node, I/O group, MDisk, MDisk group, subsystem                                                              | The peak (worst) response time among all read operations for a component over a specified time interval. For SAN Volume Controller, it represents the external response time of the MDisks.                                                      |
| Peak Back-end Write Response Time | SVC: node, I/O group, MDisk, MDisk group, subsystem                                                              | The peak (worst) response time among all write operations for a component over a specified time interval. For SAN Volume Controller, it represents the external response time of the MDisks.                                                     |
| Peak Back-end Read Queue Time     | SVC: node, I/O group, MDisk, MDisk group, subsystem                                                              | The lower bound on the peak (worst) queue time for read operations for a component over a specified time interval. The queue time is the amount of time that the read operation spent on the queue before being issued to the back-end device.   |
| Peak Back-end Write Queue Time    | SVC: node, I/O group, MDisk, MDisk group, subsystem                                                              | The lower bound on the peak (worst) queue time for write operations for a component over a specified time interval. The queue time is the amount of time that the write operation spent on the queue before being issued to the back-end device. |
| <b>Transfer sizes</b>             |                                                                                                                  |                                                                                                                                                                                                                                                  |
| Back-end Read Transfer Size       | ESS/DS6000/DS8000: rank, array, controller, subsystem<br><br>SVC: node, I/O group, MDisk, MDisk group, subsystem | Average number of KB per I/O for read operations for a component over a specified time interval.                                                                                                                                                 |
| Back-end Write Transfer Size      | ESS/DS6000/DS8000: rank, array, controller, subsystem<br><br>SVC: node, I/O group, MDisk, MDisk group, subsystem | Average number of KB per I/O for write operations for a component over a specified time interval.                                                                                                                                                |
| Overall Back-end Transfer Size    | ESS/DS6000/DS8000: rank, array, controller, subsystem<br><br>SVC: node, I/O group, MDisk, MDisk group, subsystem | Average number of KB per I/O for read and write operations for a component over a specified time interval.                                                                                                                                       |
| <b>Disk utilization</b>           |                                                                                                                  |                                                                                                                                                                                                                                                  |

| Column                      | Devices: components      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Disk Utilization Percentage | ESS/DS6000/DS8000: array | The approximate utilization percentage of a rank over a specified time interval (the average percent of time that the disks associated with the array were busy).<br><b>Note:</b> Tivoli Storage Productivity Center does not calculate a value for this column if there are multiple ranks in the extent pool where the space-efficient volumes are allocated. This column displays value of <b>N/A</b> for the reports in which it appears. However, if there is only a single rank in the extent pool, Tivoli Storage Productivity Center does calculate the value for this column regardless of the space-efficient volumes. |
| Sequential I/O Percentage   | ESS/DS6000/DS8000: array | Percentage of all I/O operations performed for an array over a specified time interval that were sequential operations.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## Frontend and fabric-based metrics

The following table contains information on frontend and fabric-based metrics.

| Column                     | Devices: components                                                                                                       | Description                                                                                                            |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------|
| <b>I/O or packet rates</b> |                                                                                                                           |                                                                                                                        |
| Port Send I/O Rate         | ESS/DS6000/DS8000: port, subsystem<br>SVC: port, node, I/O group, subsystem<br>SMI-S BSP: port<br>XIV <sup>2</sup> : port | Average number of I/O operations per second for send operations for a port over a specified time interval.             |
| Port Receive I/O Rate      | ESS/DS6000/DS8000: port, subsystem<br>SVC: port, node, I/O group, subsystem<br>SMI-S BSP: port<br>XIV <sup>2</sup> : port | Average number of I/O operations per second for receive operations for a port over a specified time interval.          |
| Total Port I/O Rate        | ESS/DS6000/DS8000: port, subsystem<br>SVC: port, node, I/O group, subsystem<br>SMI-S BSP: port<br>XIV <sup>2</sup> : port | Average number of I/O operations per second for send and receive operations for a port over a specified time interval. |
| Port Send Packet Rate      | switch port, switch                                                                                                       | Average number of packets per second for send operations for a port over a specified time interval.                    |
| Port Receive Packet Rate   | switch port, switch                                                                                                       | Average number of packets per second for receive operations for a port over a specified time interval.                 |

| Column                               | Devices: components                   | Description                                                                                                                                                              |
|--------------------------------------|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Total port Packet Rate               | switch port, switch                   | Average number of packets per second for send and receive operations for a port over a specified time interval.                                                          |
| Port to Host Send I/O Rate           | SVC: port, node, I/O group, subsystem | Average number of exchanges (I/Os) per second sent to host computers by a component over a specified time interval.                                                      |
| Port to Host Receive I/O Rate        | SVC: port, node, I/O group, subsystem | Average number of exchanges (I/Os) per second received from host computers by a component over a specified time interval.                                                |
| Total Port to Host I/O Rate          | SVC: port, node, I/O group, subsystem | Average number of exchanges (I/Os) per second transmitted between host computers and a component over a specified time interval.                                         |
| Port to Disk Send I/O Rate A         | SVC: port, node, I/O group, subsystem | Average number of exchanges (I/Os) per second sent to storage subsystems by a component over a specified time interval.                                                  |
| Port to Disk Receive I/O Rate        | SVC: port, node, I/O group, subsystem | Average number of exchanges (I/Os) per second received from storage subsystems by a component over a specified time interval.                                            |
| Total Port to Disk I/O Rate          | SVC: port, node, I/O group, subsystem | Average number of exchanges (I/Os) per second transmitted between storage subsystems and a component over a specified time interval.                                     |
| Port to Local Node Send I/O Rate     | SVC: port, node, I/O group, subsystem | Average number of exchanges (I/Os) per second sent to other nodes in the local SAN Volume Controller cluster by a component over a specified time interval.              |
| Port to Local Node Receive I/O Rate  | SVC: port, node, I/O group, subsystem | Average number of exchanges (I/Os) per second received from other nodes in the local SAN Volume Controller cluster by a component over a specified time interval.        |
| Total Port to Local Node I/O Rate    | SVC: port, node, I/O group, subsystem | Average number of exchanges (I/Os) per second transmitted between other nodes in the local SAN Volume Controller cluster and a component over a specified time interval. |
| Port to Remote Node Send I/O Rate    | SVC: port, node, I/O group, subsystem | Average number of exchanges (I/Os) per second sent to nodes in the remote SAN Volume Controller cluster by a component over a specified time interval.                   |
| Port to Remote Node Receive I/O Rate | SVC: port, node, I/O group, subsystem | Average number of exchanges (I/Os) per second received from nodes in the remote SAN Volume Controller cluster.                                                           |
| Total Port to Remote Node I/O Rate   | SVC: port, node, I/O group, subsystem | Average number of exchanges (I/Os) per second transmitted between nodes in the remote SAN Volume Controller cluster and a component over a specified time interval.      |
| Port FCP Send I/O Rate*              | ESS/DS6000/DS8000: port               | Average number of send operations per second using the FCP protocol, for a port over a specified time interval.                                                          |
| Port FCP Receive I/O Rate*           | ESS/DS6000/DS8000: port               | Average number of receive operations per second using the FCP protocol for a port over a specified time interval.                                                        |

| Column                       | Devices: components                                                                                                                                              | Description                                                                                                                                                     |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Total Port FCP I/O Rate*     | ESS/DS6000/DS8000: port                                                                                                                                          | Average number of send and receive operations per second using the FCP protocol for a port over a specified time interval.                                      |
| Port FICON Send I/O Rate*    | ESS/DS6000/DS8000: port                                                                                                                                          | Average number of send operations per second using the FICON protocol for a port over a specified time interval.                                                |
| Port FICON Receive I/O Rate* | ESS/DS6000/DS8000: port                                                                                                                                          | Average number of receive operations per second using the FICON protocol for a port over a specified time interval.                                             |
| Total Port FICON I/O Rate*   | ESS/DS6000/DS8000: port                                                                                                                                          | Average number of send and receive operations per second using the FICON protocol for a port over a specified time interval.                                    |
| Port PPRC Send I/O Rate      | ESS/DS6000/DS8000: port, subsystem                                                                                                                               | Average number of send operations per second for Peer-to-Peer Remote Copy usage for a port over a specified time interval.                                      |
| Port PPRC Receive I/O Rate   | ESS/DS6000/DS8000: port, subsystem                                                                                                                               | Average number of receive operations per second for Peer-to-Peer Remote Copy usage for a port over a specified time interval.                                   |
| Total Port PPRC I/O Rate     | ESS/DS6000/DS8000: port, subsystem                                                                                                                               | Average number of send and receive operations per second for Peer-to-Peer Remote Copy usage for a port over a specified time interval.                          |
| <b>Data rates</b>            |                                                                                                                                                                  |                                                                                                                                                                 |
| Port Send Data Rate          | ESS/DS6000/DS8000: port, subsystem<br><br>SVC: port, node, I/O group, subsystem<br><br>SMI-S BSP: port<br><br>switch port, switch<br><br>XIV <sup>2</sup> : port | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for send (read) operations for a port over a specified time interval.      |
| Port Receive Data Rate       | ESS/DS6000/DS8000: port, subsystem<br><br>SVC: port, node, I/O group, subsystem<br><br>SMI-S BSP: port<br><br>switch port, switch<br><br>XIV <sup>2</sup> : port | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for receive (write) operations for a port over a specified time interval.  |
| Total Port Data Rate         | ESS/DS6000/DS8000: port, subsystem<br><br>SVC: port, node, I/O group, subsystem<br><br>SMI-S BSP: port<br><br>switch port, switch<br><br>XIV <sup>2</sup> : port | Average number of megabytes (2 <sup>20</sup> bytes) per second that were transferred for send and receive operations for a port over a specified time interval. |

| Column                                | Devices: components                   | Description                                                                                                                                                                               |
|---------------------------------------|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port Peak Send Data Rate              | switch port                           | Peak number of megabytes (2 <sup>20</sup> bytes) per second that were sent by a port over a specified time interval                                                                       |
| Port Peak Receive Data Rate           | switch port                           | Peak number of megabytes (2 <sup>20</sup> bytes) per second that were received by a port over a specified time interval.                                                                  |
| Port to Host Send Data Rate           | SVC: port, node, I/O group, subsystem | Average number of megabytes (2 <sup>20</sup> bytes) per second sent to host computers by a component over a specified time interval.                                                      |
| Port to Host Receive Data Rate        | SVC: port, node, I/O group, subsystem | Average number of megabytes (2 <sup>20</sup> bytes) per second received from host computers by a component over a specified time interval.                                                |
| Total Port to Host Data Rate          | SVC: port, node, I/O group, subsystem | Average number of megabytes (2 <sup>20</sup> bytes) per second transmitted between host computers and a component over a specified time interval.                                         |
| Port to Disk Send Data Rate           | SVC: port, node, I/O group, subsystem | Average number of megabytes (2 <sup>20</sup> bytes) per second sent to storage subsystems by a component over a specified time interval.                                                  |
| Port to Disk Receive Data Rate        | SVC: port, node, I/O group, subsystem | Average number of megabytes (2 <sup>20</sup> bytes) per second received from storage subsystems by a component over a specified time interval.                                            |
| Total Port to Disk Data Rate          | SVC: port, node, I/O group, subsystem | Average number of megabytes (2 <sup>20</sup> bytes) per second transmitted between storage subsystems and a component over a specified time interval.                                     |
| Port to Local Node Send Data Rate     | SVC: port, node, I/O group, subsystem | Average number of megabytes (2 <sup>20</sup> bytes) per second sent to other nodes in the local SAN Volume Controller cluster by a component over a specified time interval.              |
| Port to Local nNode Receive Data Rate | SVC: port, node, I/O group, subsystem | Average number of megabytes (2 <sup>20</sup> bytes) per second received from other nodes in the local SAN Volume Controller cluster by a component over a specified time interval.        |
| Total Port to Local Node Data Rate    | SVC: port, node, I/O group, subsystem | Average number of megabytes (2 <sup>20</sup> bytes) per second transmitted between other nodes in the local SAN Volume Controller cluster and a component over a specified time interval. |
| Port to Remote Node Send Data Rate    | SVC: port, node, I/O group, subsystem | Average number of megabytes (2 <sup>20</sup> bytes) per second sent to nodes in the remote SAN Volume Controller cluster by a component over a specified time interval.                   |
| Port to Remote Node Receive Data Rate | SVC: port, node, I/O group, subsystem | Average number of megabytes (2 <sup>20</sup> bytes) per second received from nodes in the remote SAN Volume Controller cluster by a component over a specified time interval.             |
| Total Port to Remote Node Data Rate   | SVC: port, node, I/O group, subsystem | Average number of megabytes (2 <sup>20</sup> bytes) per second transmitted between nodes in the remote SAN Volume Controller cluster and a component over a specified time interval.      |
| Port FCP Send Data Rate               | ESS/DS6000/DS8000: port               | Average number of megabytes (2 <sup>20</sup> bytes) per second sent over the FCP protocol for a port over a specified time interval.                                                      |



| Column                                   | Devices: components                                           | Description                                                                                                                                                                                                                                           |
|------------------------------------------|---------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port FCP Receive Data Rate*              | ESS/DS6000/DS8000: port                                       | Average number of megabytes (2 <sup>20</sup> bytes) per second received over the FCP protocol for a port over a specified time interval.                                                                                                              |
| Total Port FCP Data Rate*                | ESS/DS6000/DS8000: port                                       | Average number of megabytes (2 <sup>20</sup> bytes) per second sent or received over the FCP protocol for a port over a specified time interval.                                                                                                      |
| Port FICON Send Data Rate*               | ESS/DS6000/DS8000: port                                       | Average number of megabytes (2 <sup>20</sup> bytes) per second sent over the FICON protocol for a port over a specified time interval.                                                                                                                |
| Port FICON Receive Data Rate*            | ESS/DS6000/DS8000: port                                       | Average number of megabytes (2 <sup>20</sup> bytes) per second received over the FICON protocol, for a port over a specified time interval.                                                                                                           |
| Total Port FICON Data Rate*              | ESS/DS6000/DS8000: port                                       | Average number of megabytes (2 <sup>20</sup> bytes) per second sent or received over the FICON protocol for a port over a specified time interval.                                                                                                    |
| Port PPRC Send Data Rate                 | ESS/DS6000/DS8000: port, subsystem                            | Average number of megabytes (2 <sup>20</sup> bytes) per second sent for Peer-to-Peer Remote Copy usage for a port over a specified time interval.                                                                                                     |
| Port PPRC Receive Data Rate              | ESS/DS6000/DS8000: port, subsystem                            | Average number of megabytes (2 <sup>20</sup> bytes) per second received for Peer-to-Peer Remote Copy usage for a port over a specified time interval.                                                                                                 |
| Total Port PPRC Data Rate                | ESS/DS6000/DS8000: port, subsystem                            | Average number of megabytes (2 <sup>20</sup> bytes) per second transferred for Peer-to-Peer Remote Copy usage for a port over a specified time interval.                                                                                              |
| <b>Response times</b>                    |                                                               |                                                                                                                                                                                                                                                       |
| Port Send Response Time                  | ESS/DS6000/DS8000: port, subsystem<br>XIV <sup>2</sup> : port | Average number of milliseconds that it took to service each send (read) operation for a port over a specified time interval.                                                                                                                          |
| Port Receive Response Time               | ESS/DS6000/DS8000: port, subsystem<br>XIV <sup>2</sup> : port | Average number of milliseconds that it took to service each receive (write) operation for a port over a specified time interval.                                                                                                                      |
| Overall Port Response Time               | ESS/DS6000/DS8000: port, subsystem<br>XIV <sup>2</sup> : port | Average number of milliseconds that it took to service each operation (send and receive) for a port over a specified time interval.                                                                                                                   |
| Port to Local Node Send Response Time    | SVC: node, I/O group, subsystem                               | Average number of milliseconds it took to service each send operation to another node in the local SAN Volume Controller cluster for a component over a specified time interval. This is the external response time of the transfers.                 |
| Port to Local Node Receive Response Time | SVC: node, I/O group, subsystem                               | Average number of milliseconds it took to service each receive operation from another node in the local SAN Volume Controller cluster for a component over a specified time interval. This is the external response time of the transfers.            |
| Total Port to Local Node Response Time   | SVC: node, I/O group, subsystem                               | Average number of milliseconds it took to service each send or receive operation between another node in the local SAN Volume Controller cluster and a component over a specified time interval. This is the external response time of the transfers. |

| Column                                    | Devices: components             | Description                                                                                                                                                                                                                                      |
|-------------------------------------------|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port to Local Node Send Queued Time       | SVC: node, I/O group, subsystem | Average number of milliseconds that each send operation issued to another node in the local SAN Volume Controller cluster spent on the queue before being issued for a component over a specified time interval.                                 |
| Port to Local Node Receive Queued Time    | SVC: node, I/O group, subsystem | Average number of milliseconds that each receive operation from another node in the local SAN Volume Controller cluster spent on the queue before being issued for a component over a specified time interval.                                   |
| Total Port to Local Node Queued Time      | SVC: node, I/O group, subsystem | Average number of milliseconds that each operation issued to another node in the local SAN Volume Controller cluster spent on the queue before being issued for a component over a specified time interval.                                      |
| Port to Remote Node Send Response Time    | SVC: node, I/O group, subsystem | Average number of milliseconds it took to service each send operation to a node in the remote SAN Volume Controller cluster for a component over a specified time interval. This is the external response time of the transfers.                 |
| Port to Remote Node Receive Response Time | SVC: node, I/O group, subsystem | Average number of milliseconds it took to service each receive operation from a node in the remote SAN Volume Controller cluster for a component over a specified time interval. This is the external response time of the transfers.            |
| Total Port to Remote Node Response Time   | SVC: node, I/O group, subsystem | Average number of milliseconds it took to service each send or receive operation between a node in the remote SAN Volume Controller cluster and a component over a specified time interval. This is the external response time of the transfers. |
| Port to Remote Node Send Queued Time      | SVC: node, I/O group, subsystem | Average number of milliseconds that each send operation issued to a node in the remote SAN Volume Controller cluster spent on the queue before being issued for a component over a specified time interval.                                      |
| Port to Remote Node Receive Queued Time   | SVC: node, I/O group, subsystem | Average number of milliseconds that each receive operation from a node in the remote SAN Volume Controller cluster spent on the queue before being issued for a component over a specified time interval.                                        |
| Total Port to Remote Node Queued Time     | SVC: node, I/O group, subsystem | Average number of milliseconds that each operation issued to a node in the remote SAN Volume Controller cluster spent on the queue before being issued for a component over a specified time interval.                                           |
| Port FCP Send Response Time*              | ESS/DS6000/DS8000: port         | Average number of milliseconds it took to service all send operations over the FCP protocol for a port over a specified time interval.                                                                                                           |

| Column                            | Devices: components                                   | Description                                                                                                                                            |
|-----------------------------------|-------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| Port FCP Receive Response Time*   | ESS/DS6000/DS8000: port                               | Average number of milliseconds it took to service all receive operations over the FCP protocol for a port over a specified time interval.              |
| Overall Port FCP Response Time*   | ESS/DS6000/DS8000: port                               | Average number of milliseconds it took to service all I/O operations over the FCP protocol for a port over a specified time interval.                  |
| Port FICON Send Response Time*    | ESS/DS6000/DS8000: port                               | Average number of milliseconds it took to service all send operations over the FICON protocol for a port over a specified time interval.               |
| Port FICON Receive Response Time* | ESS/DS6000/DS8000: port                               | Average number of milliseconds it took to service all receive operations over the FICON protocol for a port over a specified time interval.            |
| Overall Port FICON Response Time* | ESS/DS6000/DS8000: port                               | Average number of milliseconds it took to service all I/O operations over the FICON protocol for a port over a specified time interval.                |
| Port PPRC Send Response Time      | ESS/DS6000/DS8000: port, subsystem                    | Average number of milliseconds it took to service all send operations for Peer-to-Peer Remote Copy usage for a port over a specified time interval.    |
| Port PPRC Receive Response Time   | ESS/DS6000/DS8000: port, subsystem                    | Average number of milliseconds it took to service all receive operations for Peer-to-Peer Remote Copy usage for a port over a specified time interval. |
| Overall Port PPRC Response Time   | ESS/DS6000/DS8000: port, subsystem                    | Average number of milliseconds it took to service all I/O operations for Peer-to-Peer Remote Copy usage for a port over a specified time interval.     |
| <b>Transfer sizes</b>             |                                                       |                                                                                                                                                        |
| Port Send Transfer Size           | ESS/DS6000/DS8000: port, subsystem<br>SMI-S BSP: port | Average number of KB sent per I/O by a port over a specified time interval.                                                                            |
| Port Receive Transfer Size        | ESS/DS6000/DS8000: port, subsystem<br>SMI-S BSP: port | Average number of KB received per I/O by a port over a specified time interval.                                                                        |
| Overall Port Transfer Size        | ESS/DS6000/DS8000: port, subsystem<br>SMI-S BSP: port | Average number of KB transferred per I/O by a port over a specified time interval.                                                                     |
| Port Send Packet Size             | switch port, switch                                   | Average number of KB sent per packet by a port over a specified time interval.                                                                         |
| Port Receive Packet Size          | switch port, switch                                   |                                                                                                                                                        |
| Overall Port Packet Size          | switch port, switch                                   | Average number of KB received per packet by a port over a specified time interval.                                                                     |
| <b>Special computed values</b>    |                                                       |                                                                                                                                                        |

| Column                              | Devices: components                                                                           | Description                                                                                                                       |
|-------------------------------------|-----------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| Port Send Utilization Percentage    | ESS/DS6000/DS8000: port                                                                       | Average amount of time that the port was busy sending data over a specified time interval.                                        |
| Port Receive Utilization Percentage | ESS/DS6000/DS8000: port                                                                       | Average amount of time that the port was busy receiving data over a specified time interval.                                      |
| Overall Port Utilization Percentage | ESS/DS6000/DS8000: port                                                                       | Average amount of time that the port was busy sending or receiving data over a specified time interval.                           |
| Port Send Bandwidth Percentage      | ESS/DS8000: port<br>SVC: port<br>switch, port<br>XIV <sup>2</sup> : port                      | The approximate bandwidth utilization percentage for send operations by a port based on its current negotiated speed.             |
| Port Receive Bandwidth Percentage   | ESS/DS8000: port<br>SVC: port<br>switch, port<br>XIV <sup>2</sup> : port                      | The approximate bandwidth utilization percentage for receive operations by this port, based on its current negotiated speed.      |
| Overall Port Bandwidth Percentage   | ESS/DS8000: port<br>SVC: port<br>switch, port<br>XIV <sup>2</sup> : port                      | The approximate bandwidth utilization percentage for send and receive operations by this port.                                    |
| <b>Error rates</b>                  |                                                                                               |                                                                                                                                   |
| Error Frame Rate                    | switch port, switch<br>DS8000: port,<br>subsystem                                             | The number of frames per second that were received in error by a port over a specified time interval.                             |
| Dumped Frame Rate                   | switch port, switch                                                                           | The number of frames per second that were lost due to a lack of available host buffers for a port over a specified time interval. |
| Link Failure Rate                   | switch port, switch<br>DS8000: port,<br>subsystem<br>SVC: port, node, I/O<br>group, subsystem | The number of link errors per second that were experienced by a port over a specified time interval.                              |
| Loss of Sync Rate                   | switch port, switch<br>DS8000: port,<br>subsystem<br>SVC: port, node, I/O<br>group, subsystem | The average number of times per second that synchronization was lost for a component over a specified time interval.              |

| Column                                 | Devices: components                                                                             | Description                                                                                                                                                                                                                                                  |
|----------------------------------------|-------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Loss of Signal Rate                    | switch port, switch<br><br>DS8000: port, subsystem<br><br>SVC: port, node, I/O group, subsystem | The average number of times per second that the signal was lost for a component over a specified time interval.                                                                                                                                              |
| CRC Error Rate                         | switch port, switch<br><br>DS8000: port, subsystem<br><br>SVC: port, node, I/O group, subsystem | The average number of frames received per second in which the CRC in the frame did not match the CRC computed by the receiver for a component over a specified time interval.                                                                                |
| Short Frame Rate                       | switch port, switch                                                                             | The average number of frames received per second that were shorter than 28 octets (24 header + 4 CRC) not including any SOF/EOF bytes for a component over a specified time interval.                                                                        |
| Long Frame Rate                        | switch port, switch                                                                             | The average number of frames received per second that were longer than 2140 octets (24 header + 4 CRC + 2112 data) not including any SOF/EOF bytes for a component over a specified time interval.                                                           |
| Encoding Disparity Error Rate          | switch port, switch                                                                             | The average number of disparity errors received per second for a component over a specified time interval.                                                                                                                                                   |
| Discarded Class3 Frame Rate            | switch port, switch                                                                             | The average number of class-3 frames per second that were discarded by a component over a specified time interval.                                                                                                                                           |
| F-BSY Frame Rate                       | switch port, switch                                                                             | The average number of F-BSY frames per second that were generated by a component over a specified time interval.                                                                                                                                             |
| F-RJT Frame Rate                       | switch port, switch                                                                             | The average number of F-RJT frames per second that were generated by a component over a specified time interval.                                                                                                                                             |
| Primitive Sequence Protocol Error Rate | switch port, switch<br><br>DS8000: port, subsystem<br><br>SVC: port, node, I/O group, subsystem | The average number of primitive sequence protocol errors detected for a component over a specified time interval.                                                                                                                                            |
| Invalid Transmission Word Rate         | switch port, switch<br><br>DS8000: port, subsystem<br><br>SVC: port, node, I/O group, subsystem | The average number of transmission words per second that had an 8b10 code violation in one or more characters; had a K28.5 in its second, third, or fourth character positions; and/or was an ordered set that had an incorrect Beginning Running Disparity. |
| Zero Buffer-Buffer Credit Timer        | SVC: port, node, I/O group, subsystem                                                           | The number of microseconds that the port has been unable to send frames due to lack of buffer credit since the last node reset.                                                                                                                              |

| Column                                                                                                                                           | Devices: components                            | Description                                                                                                                                                                               |
|--------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Link Recovery (LR) Sent Rate                                                                                                                     | switch port, switch<br>DS8000: port, subsystem | The average number of times per second that a port has transitioned from an active (AC) state to a Link Recovery (LR1) state over a specified time interval.                              |
| Link Recovery (LR) Received Rate                                                                                                                 | switch port, switch<br>DS8000: port, subsystem | The average number of times per second a port has transitioned from an active (AC) state to a Link Recovery (LR2) state over a specified time interval                                    |
| Out of Order Data Rate                                                                                                                           | DS8000: port, subsystem                        | The average number of times per second that an out of order frame was detected for a port over a specified time interval.                                                                 |
| Out of Order ACK Rate                                                                                                                            | DS8000: port, subsystem                        | The average number of times per second that an out of order ACK frame was detected for a port over a specified time interval.                                                             |
| Duplicate Frame Rate                                                                                                                             | DS8000: port, subsystem                        | The average number of times per second that a frame was received that has been detected as previously processed for a port over a specified time interval.                                |
| Invalid Relative Offset Rate                                                                                                                     | DS8000: port, subsystem                        | The average number of times per second that a frame was received with invalid relative offset in the frame header for a port over a specific time interval.                               |
| Sequence Timeout Rate                                                                                                                            | DS8000: port, subsystem                        | The average number of times per second the port has detected a timeout condition on receiving sequence initiative for a fiber channel exchange for a port over a specified time interval. |
| <b>Note:</b> * The value N/A is displayed for this metric if you set the <b>Summation Level</b> to hourly or daily before generating the report. |                                                |                                                                                                                                                                                           |

---

## Protocols and standards

This section provides an overview of the protocols and standards that are used within IBM Tivoli Storage Productivity Center.

### Web Based Enterprise Management

Web Based Enterprise Management (WBEM) is an initiative of the Distributed Management Task Force (DMTF) with the objective to enable the management of complex IT environments. It defines a set of management and internet standard technologies in order to unify the management of complex IT environments.

The WBEM initiative is composed of three main conceptual elements:

#### Common Interface Model (CIM)

CIM is a formal object-oriented modeling language that is used to describe the management aspects of systems.

#### xmlCIM

This is the grammar to describe CIM declarations and messages used by the CIM protocol.

#### Hypertext Transfer Protocol (HTTP)

HTTP is used as a way to enable communication between a management application and a device that both use CIM.

The WBEM architecture defines the following elements:

**CIM Client**

The CIM Client is a management application like IBM Tivoli Storage Productivity Center that uses CIM to manage devices. A CIM Client can reside anywhere in the network, because it uses HTTP to talk to CIM Object Managers and Agents.

**CIM Managed Object**

A Managed Object is a hardware or software component that can be managed by a management application by using CIM.

**CIM Agent**

A CIM Object Manager that includes the provider service for a limited set of resources. An agent may be embedded or hosted and can be an aggregator for multiple devices.

**CIM Provider**

A CIM Provider is the element that translates CIM calls to the device-specific commands. A provider is always closely linked to a CIM.

**CIM Object Manager (CIMOM)**

The central component of the CIM Server responsible for the communication between the CIM server components.

**CIM Server**

A server that receives and processes CIM Operation Message Requests and issues CIM Operation Message Responses.

## **Storage Management Initiative Specification**

The Storage Networking Industry Association (SNIA) defines a standard that is used within IBM Tivoli Storage Productivity Center to create and develop a universal open interface for managing storage devices including storage networks.

Information on SMI-S can be found at this Web site:

<http://www.snia.org>

SNIA has fully adopted and enhanced the Common Information Model (CIM) standard for storage management in its Storage Management Initiative - Specification (SMI-S). SMI-S was launched to create and develop a universal open interface for managing storage devices including storage networks. SMI-S provides:

- A comprehensive specification for the management of heterogeneous storage and storage area networks (SANs).
- The information available to a WBEM client from an SMI-S compliant CIM server (provider).
- Profiles organized by:
  - Storage
  - Fabric
  - Host
  - Common profiles and subprofiles
- An object-oriented CIM and XML-based interface for managing SAN devices, services, and fabrics.
- An initial discovery, which is SLP based.

The idea behind SMI-S is to standardize the management interfaces so that management applications can utilize these and provide cross-device management. This means that a newly introduced device can be immediately managed as it will conform to the standards.

The models and protocols in the SMI-S implementation are platform-independent, enabling application development for any platform, and enabling them to run on different platforms. The SNIA will also provide interoperability tests which will help vendors test their applications and devices if they conform to the standard.

## Service Location Protocol

The Service Location Protocol (SLP) is an Internet Engineering Task Force (IETF) standard. SLP provides a scalable framework for the discovery and selection of network services.

The Internet Engineering Task Force (IETF) is a large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. The IETF includes formal standards for SNMP and MIBs. More information on IETF can be found at this Web site:

<http://www.ietf.org>

SLP enables the discovery and selection of generic services, which could range in function from hardware services such as those for printers or fax machines, to software services such as those for file servers, email servers, Web servers, databases, or any other possible services that are accessible through an IP network.

Traditionally, to use a particular service, a user or client application needed to supply the host name or network IP address of that service. With SLP, however, the user or client no longer needs to know individual host names or IP addresses. Instead, the user or client can search the network for the desired service type and an optional set of qualifying attributes.

For example, a user could specify to search for all available printers that support Postscript. Based on the given service type (printers), and the given attributes (Postscript). SLP searches the user's network for any matching services, and returns the discovered list to the user.

## Simple Network Management Protocol

The Simple Network Management Protocol (SNMP) is an Internet Engineering Task Force (IETF) protocol for monitoring and managing systems and devices in a network. Functions supported by the SNMP protocol are the request and retrieval of data, the setting or writing of data, and traps that signal the occurrence of events.

SNMP is a method that enables a management application to query information from a managed device. The managed device has software running that sends and receives the SNMP information. This software module is usually called the SNMP agent.

An SNMP manager can read information from an SNMP agent in order to monitor a device. Therefore, the device needs to be polled on an interval bases. The SNMP manager can also change the configuration of a device, by setting certain values to



corresponding variables. A device can also be set up to send a notification to the SNMP manager (this is called a trap) to asynchronously inform this SNMP manager of a status change.

Depending on the existing environment and organization it is very likely that your environment already has an SNMP management application in place. IBM Tivoli Storage Productivity Center can be set up to send traps.

SNMP uses a hierarchical structured Management Information Base (MIB) to define the meaning and the type of a particular value. A MIB defines managed objects that describe the behavior of the SNMP entity, which can be anything from an IP router to a storage subsystem. The information is organized in a tree structure.

For users planning to make use of the Tivoli Storage Productivity Center SNMP trap alert notification capabilities, an SNMP MIB is included in the server installation.

The MIB is provided for use by your SNMP management console software. Most SNMP management products provide a program called a MIB compiler that can be used to import MIBs. This will allow you to view Tivoli Storage Productivity Center generated SNMP traps from within your management console software. Refer to your management console software documentation for instructions on how to compile or import a third party MIB.

Tivoli Storage Productivity Center uses port 162 to listen for SNMP traps. This is the default port. For switches, you must configure the switch to send SNMP traps to the Device server IP address. If you need to change the default port number, use the `setdscfg` command. The attribute to set is `SNMPTrapPort`.

System administrators must set up their SNMP trap ringer with the provided MIB files in order to receive SNMP traps from Tivoli Storage Productivity Center. These files are located in the following directories on the product installation CD:

For Data server:

```
data\snmp\tivoliSRM.mib
```

For Device server:

```
device\snmp\fabric.mib
```

## **IBM Tivoli Storage Productivity Center Universal Agent**

IBM Tivoli Storage Productivity Center Universal Agent collects information about the Tivoli Storage Productivity Center Health, Data Server, and Data Server services information, Device information, alert information, and job information.

### **Tivoli Storage Productivity Center Health, Data Server and Data Server services Information**

This will include:

- Data server status
- Device server status
- Services status
- CIMOM connection status
- CIMOM information such as last Discovery, Managed devices
- Agent connection status

- Agent information such as last Discovery and Probe, Managed devices
- Equivalent information about other Tivoli Storage Productivity Center servers, VM Ware

### **Alert Information**

- Amount of all alerts
- Alerts per component such as Computer, Data, Disk, Fabric etc.

### **Job Information**

- Amount of Jobs, such as Discovery, Probe, Scans, PM
- Job Status and Details, For example: Start Time, Finish Time, Status, Log File Name
- Scheduled Jobs and Details, For example: Intervals, Creator, Name

For more information, see the disk1 image in the tools directory.

## **Fibre Channel Methodologies of Interconnects**

IBM Tivoli Storage Productivity Center supports the ANSI T11 Fibre Channel FC-MI (Fibre Channel Methodologies of Interconnects) for the automated discovery of FC SAN assets and topology.

ANSI T11 Fibre Channel FC-MI includes the following for the automated discovery of FC SAN assets and topology:

- Hosts (HBAs)
- FC interconnects
- FC storage devices

The T11 FC-MI also includes the following:

- FC-GS-3/4 (discovery, zoning, and so forth)
- RNID (advanced device recognition)
- Platform registration (device recognition and launch)
- Common HBA API (fabric and storage views)
- Name server (connectivity)
- Management server (SAN connectivity and topology)
- RSCN (advanced event detection)
- SCSI queries (storage views, volume information, and so forth)
- SNMP Fabric Element (FE) MIB
- SNMP FC Management MIB (discovery, performance statistics, and so forth)
- SNMP alerts

---

## Appendix B. Accessibility features for Tivoli Storage Productivity Center

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

### Accessibility features

The following list includes the major accessibility features in IBM Tivoli Storage Productivity Center:

- Tivoli Storage Productivity Center functions are available using the keyboard for navigation instead of the mouse. You can use keys or key combinations to perform operations that can also be done using a mouse. However, you must use the mouse to navigate the Topology Viewer and report graphs. Standard operating system keystrokes are used for standard operating system operations.
- You can use screen readers to read the user interface.
- The user interface communicates all information independently of color.
- The *IBM Tivoli Storage Productivity Center Information Center*, and its related publications are accessibility-enabled and include the following accessibility features:
  - The information center is provided in XHTML 1.0 format, which is viewable in most Web browsers. XHTML allows you to view documentation according to the display preferences set in your browser. With XHTML, you can use screen readers and other assistive technologies.
  - All documentation for Tivoli Storage Productivity Center is available in Adobe Portable Document Format (PDF) using the Adobe Acrobat Reader. You can access the PDFs from the Printable PDFs topic in the information center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp>.
  - All images are provided with alternative text, so that users with vision impairments can understand the contents of the images.

### Keyboard navigation

This product uses standard Microsoft Windows navigation keys.

### Displaying the interface

Click **Preferences > Look and Feel** to select how to display the Tivoli Storage Productivity Center user interface. To do so, complete the following steps:

1. Start the Tivoli Storage Productivity Center user interface.
2. Click **Preferences > Look and Feel** to change the visual appearance of the user interface to best suit your visual needs:
  - Windows Classic
  - Windows
  - CDE/Motif
  - Metal

## **Limitations for sight-impaired users on IBM Tivoli Storage Productivity Center for Replication**

There are some limitations in the product graphical user interface (GUI) that require special attention before you start using Tivoli Storage Productivity Center for Replication with a screen reader:

**Dojo components are not read by all screen readers.**

The JAWS screen reader does not read some Dojo components on Internet Explorer 7. Use the command-line interface instead of the GUI with JAWS on Internet Explorer 7.

### **IBM and accessibility**

See the IBM Human Ability and Accessibility Center Web site for more information about the commitment that IBM has to accessibility.

---

## Appendix C. Accessibility features for Tivoli Storage Productivity Center for Replication

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

The following list includes the major accessibility features in Tivoli Storage Productivity Center for Replication:

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Keys that are discernible by touch but do not activate just by touching them
- Industry-standard devices for ports and connectors
- The attachment of alternative input and output devices

See the IBM Human Ability and Accessibility Center website at [www.ibm.com/able](http://www.ibm.com/able) for more information about the commitment that IBM has for accessibility.

### Accessibility and keyboard shortcuts in the information center

Accessibility features help users with physical disabilities, such as restricted mobility or limited vision, to use software products successfully. Using the major accessibility features in this product, users can perform these tasks:

- Use assistive technologies, such as screen-reader software and digital speech synthesizer, to hear what is displayed on the screen. Consult the product documentation of the assistive technology for details on using those technologies with this product.
- Operate specific or equivalent features by using only the keyboard.
- Magnify what is displayed on the screen.

In addition, the documentation was modified to include the following features to aid accessibility:

- All documentation is available in HTML formats to give the maximum opportunity for users to apply screen-reader software technology.
- All images in the documentation are provided with alternative text so that users with vision impairments can understand the contents of the images.

Use the following key combinations to navigate the interface by keyboard:

- To go directly to the Topic pane (the right side), press Alt+K, and then press Tab.
- In the Topic pane, to go to the next link, press Tab.
- To go directly to the Search Results view in the left side, press Alt+R, and then press Enter or Up Arrow to enter the view.
- To go directly to the Navigation (Table of Contents) view in the left side, press Alt+C, and then press Enter or Up Arrow to enter the view.
- To expand and collapse a node in the navigation tree, press the Right and Left Arrow.
- To move to the next topic node, press the Down Arrow or Tab.
- To move to the previous topic node, press the Up Arrow or Shift+Tab.

- To go to the next link, button, or topic node from inside on of the views, press Tab.
- To scroll all the way up or down in a pane, press Home or End.
- To go back, press Alt+Left Arrow; to go forward, press Alt+Right Arrow.
- To go to the next pane, press F6.
- To move to the previous pane, press Shift+F6.
- To print the active pane, press Ctrl+P.

## **Related accessibility information for sight-impaired users**

The following list contains hints and tips that can help you more fully use the graphical user interface:

### **Drop-down lists are positioned directly above or before the radio button that activates it.**

If you use a screen reader, you should be aware that there are radio buttons to activate drop-down lists for several GUI pages. The way to activate the drop-down list is by selecting the associated radio button. The drop-down list is positioned directly above or before the radio button that activates it. When you use a screen reader that processes the fields and controls of a page sequentially, you might select the radio button, but not know that the associated drop-down list has been activated. The screen reader processes inactive drop-down lists first, and then processes the next radio button. The drop-down list is activated if you select the radio button.

On the following pages, keep in mind that radio buttons activate a drop-down list:

- Administration
- ESS/DS Paths
- Sessions
- Session Details
- Storage Systems

### **Tables are best understood by reviewing the surrounding text and the table row and column number of the table.**

On some graphical user pages, tables use the header or row ID attributes when reading a single cell. The screen reader reads the table row and column number, along with cell data. Therefore, you can infer the column header and row ID.

### **Experiment with and fine-tune the way your screen reader pronounces some of the product abbreviations.**

Your screen reader might pronounce abbreviations as if they were words. For example, the common abbreviation for Enterprise Storage Server is ESS. Your screen reader might read ESS as the word "ess". With some screen readers you can hear alternate pronunciations. If you frequently use the software you might prefer to fine-tune such associations in your settings. When an association is created, the screen reader can recognize the abbreviation as a word. If you can add dictionary words with your screen reader, replace the capitalized character sequence with the sequence E space S space S.

Typically, this abbreviation is used in the combination form of ESS/DS. This term refers to the Enterprise Storage Server 800, the DS6000, or the DS8000.

**Some decorative artifacts might persist if the cascading style sheet is disabled.**

Enable cascading style sheets when possible; otherwise, some decorative elements might persist in the Firefox and Internet Explorer GUIs. These artifacts do not affect performance. If they become too distracting, consider using the command-line interface instead.

**For efficiency, confirmation dialogs place initial focus on the Yes button.**

When a confirmation dialog box is displayed, focus is given to the **Yes** button. Therefore, the screen reader reads “Yes” but does not read the confirmation text. The software processes the information in this way when you do the following types of tasks:

- Perform an action on a session
- Remove a connection to a storage system
- Click the **About link**
- Create a high-availability connection

To read the confirmation text before clicking the **Yes**, **No**, or **OK** button, view the previous heading before the button.

**Dojo components are not read by all screen readers.**

The Job Access for Windows and Speech (JAWS) screen reader does not read some Dojo components on Internet Explorer 7. Use the command-line interface instead of the GUI with JAWS on Internet Explorer 7.

**Firefox is the preferred browser for use with a screen reader.**

Use Firefox as the screen reader because other browsers might not fully expose assistive technology content to the screen reader.





---

## Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan, Ltd.  
1623-14, Shimotsuruma, Yamato-shi  
Kanagawa 242-8502 Japan*

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:**

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation  
2Z4A/101  
11400 Burnet Road  
Austin, TX 78758  
U.S.A*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>

---

## Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

Adobe and PostScript are registered trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Itanium is a registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.



---

## Glossary

This glossary includes terms and definitions for IBM Tivoli Storage Productivity Center.

The following cross-references are used in this glossary:

- *See* refers the reader from a term to a preferred synonym, or from an acronym or abbreviation to the defined full form.
- *See also* refers the reader to a related or contrasting term.

To view glossaries for other IBM products, go to: <http://www.ibm.com/software/globalization/terminology/>.

### **active management server**

A management server from which the storage environment can be monitored and managed. The active management server replicates its database to the standby server.

### **business continuity**

The capability of a business to withstand outages and to operate mission-critical services normally and without interruption in accordance with predefined service-level agreements.

**CIM** See *Common Information Model*.

### **CIM agent**

The code that consists of common building blocks that can be used instead of proprietary software or device-specific programming interfaces to manage devices that are compliant with the Common Information Model (CIM).

**CKD** See *count key data*.

### **cluster**

1. In SAN Volume Controller or Storwize V7000, a pair of nodes that provides a single configuration and service interface.
2. In IBM System Storage™ DS8000®, a partition capable of performing all DS8000 series functions. With two clusters in the DS8000 storage unit, any operational cluster can take over the processing of a failing cluster.

### **Common Information Model (CIM)**

An implementation-neutral, object-oriented schema for describing

network management information. The Distributed Management Task Force (DMTF) develops and maintains CIM specifications.

### **community name**

The part of a Simple Network Management Protocol (SNMP) message that represents a password-like name and that is used to authenticate the SNMP message.

### **consistency group**

A group of copy relationships between virtual volumes or data sets that are maintained with the same time reference so that all copies are consistent in time.

### **copy set**

The source and target volumes in a session that contain copies of the same data. A session can contain one or more copy sets.

### **count key data (CKD)**

In mainframe computing, a data-record format employing self-defining record formats in which each record is represented by up to three fields: a count field identifying the record and specifying its format, an optional key field that can be used to identify the data area contents, and an optional data field that typically contains the user data.

### **data collection**

The process of obtaining performance and availability monitoring data and providing that data to a metric evaluator. Examples of data collectors include Domain Name System (DNS) probes, web page analyzers, or database analyzers. See also *discovery*.

### **discovery**

The process of finding resources within an enterprise, including finding the new location of monitored resources that were moved. See also *schedule*.

### **enterprise repository**

A component of the Data server that

records and stores all information about the monitored computers' storage assets and their usage over time. The repository is organized into relational database tables and is accessed by Data server using Java Database Connectivity (JDBC).

**extended count key data (ECKD)**

An extension of the count key data (CKD) architecture.

**FlashCopy**

In storage systems, a feature that creates a point-in-time copy of a volume or set of volumes on a specified target volume.

**flush-through mode**

A process in which data is written to a storage device at the same time as the data is cached.

**freeze** An operation in which a storage system blocks I/O from the host system to the affected volumes on the primary site. A freeze operation stops mirroring between the primary and secondary volumes to ensure data consistency at the secondary site. See also *thaw*.

**GiB** See *gibibyte*.

**gibibyte (GiB)**

A base-2 unit of measurement that is equal to 1,073,741,824 bytes (2 to the 30th power). See also *gigabyte*.

**gigabyte (GB)**

In decimal notation, 1,073,741,824 bytes, when referring to memory capacity; in all other cases, it is defined as 1,000,000,000.

**Global Copy**

A non-synchronous long-distance copy option for data migration and backup. See also *remote mirror and copy* and *Metro Mirror*.

**Global Mirror**

An optional capability of the remote mirror and copy feature that provides a two-site extended-distance remote copy. Data that is written by the host to the storage unit at the local site is automatically maintained at the remote site. See also *Metro Mirror* and *remote mirror and copy*.

**Globally Unique Identifier (GUID)**

An algorithmically determined number that uniquely identifies an entity within a system.

**grain size**

The unit size for allocating space on thin-provisioned volumes, such as 32, 64, 128, and 256 kibibyte. The grain size is defined when a volume is created.

**GUID** See *Globally Unique Identifier*.

**Hardware Management Console (HMC)**

In a system storage environment, a system that acts as the focal point for configuration, management of Copy Services functions, and maintenance.

**heat map**

A color-coded data chart where colors are used to differentiate values in a data set.

**high availability**

Pertaining to a clustered system that is reconfigured when node or daemon failures occur, so that workloads can be redistributed to the remaining nodes in the cluster.

**high-availability solution**

A combination of hardware, software, and services that fully automates the recovery process and does not disrupt user activity. High-availability solutions must provide an immediate recovery point with a fast recovery time.

**HMC** See *Hardware Management Console*.

**host volume**

A volume that represents the volume functional role from an application point of view. The host volume can be connected to a host or server. It receives read, write, and update application I/O, depending on the site to which the application is writing.

**in-band discovery**

The process of discovering information about the storage area network (SAN), including topology and attribute data, through the Fibre Channel data paths. Contrast with *out-of-band discovery*.

**job** A unit of work that monitors, manages, or collects data about a storage resource.

**job scheduler**

A component of the Data server that deploys all monitoring activities.

**journal volume**

A volume that holds a consistent copy of data until a new consistent copy is

formed. The journal volume restores the last consistent point during a recovery.

**KiB** See *kibibyte*.

**kibibyte (KiB)**

A base-2 unit of measurement that is equal to 1,024 bytes (2 to the 10th power). See also *kilobyte*.

**kilobyte**

For processor storage, real and virtual storage, and channel volume, 2 to the power of 10 or 1,024 bytes. For disk storage capacity and communications volume, 1,000 bytes.

**logical unit number (LUN)**

In the Small Computer System Interface (SCSI) standard, a unique identifier that is used to differentiate devices, each of which is a logical unit (LU).

**LUN** See *logical unit number*.

**Management Information Base (MIB)**

In the Simple Network Management Protocol (SNMP), a database of objects that can be queried or set by a network management system.

**management server**

A system that provides a central point of control for managing data replication.

**management server relationship**

A connection between two replication servers, where one server acts as the active server and replicates the data that is necessary for the standby server to take control of the replication environment.

**mebibyte (MiB)**

A base-2 unit of measurement that is equal to 1,048,576 bytes (2 to the 20th power). See also *megabyte*.

**megabyte (MB)**

For processor storage, real and virtual storage, and channel volume, 2 to the 20th power or 1,048,576 bytes. For disk storage capacity and communications volume, 1,000,000 bytes.

**Metro Mirror**

A method of synchronous remote replication that operates over distances up to 300 kilometers. This method constantly updates a secondary copy of a volume to

match changes that are made to a source volume. See also *remote mirror and copy* and *Global Copy*.

**Metro Global Mirror**

A three-site, high-availability, disaster recovery solution. Metro Global Mirror uses synchronous replication to mirror data between a local site and an intermediate site, and asynchronous replication to mirror data from an intermediate site to a remote site.

**MiB** See *mebibyte*.

**MIB** See *Management Information Base*.

**native interface**

An interface that is specific to a system or subsystem.

**out-of-band discovery**

The process of discovering storage area network (SAN) information, including topology and device data, without using the fibre-channel data paths. A common mechanism for out-of-band discovery is the use of Simple Network Management Protocol (SNMP) Management Information Base (MIB) queries, which are invoked over an IP network. Contrast with *in-band discovery*.

**ping** A job that tracks the availability of assets and that is performed by an agent. Several ping jobs can be used to monitor the availability of any computer or subset of computers in the network. See also *schedule*.

**pool** A storage pool or pool is a grouping of storage space that consists of volumes, logical unit numbers (LUNs), or addresses that share a common set of administrative characteristics.

**practice volume**

A volume that can be used to test disaster-recovery actions while maintaining disaster-recovery capability.

**primordial pool**

Storage capacity that is deallocated on a storage device. Storage pools are created by allocating storage capacity from primordial pools.

**probe** A job that itemizes and creates an inventory of assets, such as computers,

controllers, disk drives, file systems, and logical units. See also *schedule*.

**recovery point objective**

The maximum amount of data loss that can be tolerated during a service interruption.

**remote mirror and copy**

A feature of a storage server that constantly updates a secondary copy of a logical volume to match changes made to a primary logical volume. The primary and secondary volumes can be on the same storage server or on separate storage servers. See also *Metro Mirror* and *Global Mirror*.

**resource**

In a storage environment, an entity that is monitored. Resources can include fabrics, switches, computers, storage systems, and tape libraries.

**role**

A job function that identifies the tasks that a user can perform and the resources to which a user has access. A user can be assigned one or more roles.

**role pair**

The association of two volume roles in a session that take part in a copy relationship. For example, in a Metro Mirror session, the role pair can be the association between host volumes at the primary site and host volumes at the secondary site (H1-H2).

**scan job**

A job that monitors the usage and consumption of storage. See also *schedule*.

**schedule**

A monitoring task that collects data about storage resources. Monitoring tasks include scans, probes, pings, discovery, and performance monitors. For example, define a probe schedule to collect data about storage systems; define a scan schedule to collect file-level information about hosts.

**session**

A collection of source and target volumes that are managed to create consistent copies of data. The type of data replication that is associated with the session determines the actions that can be conducted for the volumes.

**site awareness**

The association of a location with each storage system in a session. Site awareness ensures that a volume can only be selected for a session if it matches the location of the site. Site awareness helps to prevent both reversing a hardware relationship and selecting volumes at the wrong location.

**SMI-S** See *Storage Management Initiative Specification*.

**standby management server**

A management server that is a backup for the active server. The replication environment cannot be monitored or managed from the standby server.

**Storage Management Initiative Specification (SMI-S)**

A design specification developed by the Storage Networking Industry Association (SNIA) that specifies a secure and reliable interface with which storage management systems (SMSs) can identify, classify, monitor, and control physical and logical resources in a storage area network (SAN). The interface integrates the various devices to be managed in a SAN and the tools used to manage them.

**storage pod**

A logical entity of components of a system that consists of two storage nodes, and one or two storage subsystems directly connected with these storage nodes.

**storage resource group**

A named collection of logically related resources that are monitored by Tivoli Storage Productivity Center. Monitored resources can include fabrics, switches, computers, storage systems, tape libraries, and other storage resource groups.

**target volume**

A volume that receives data from a host volume or another intermediate volume.

**thaw**

An operation in which a storage system releases the block of the I/O from the host system to the affected volumes on the primary site. A thaw operation can occur after a freeze operation ends and consistency is formed on the secondary site. See also *freeze*.



**track space-efficient (TSE) volume**

A volume in which storage space is allocated on an as-needed basis by using space on the target volume only when tracks are copied from the source volume to the target volume.

**virtualization**

The substitution of virtual resources for actual resources, where the virtual resources have the same functions and external interfaces as their counterparts, but differ in attributes, such as size, performance, and cost. Virtualization is commonly applied to physical hardware resources by combining multiple physical resources into shared pools from which users receive virtual resources.

**virtual storage area network (VSAN)**

A fabric within the storage area network (SAN).

**VSAN** See *virtual storage area network*.

**zone** A logical grouping of switches, switch ports, and their attached devices in a fabric.

**zone alias**

A name that is given to a collection of one or more zone members that you want to manage together.

**zone set**

A group of zones that function together on a fabric.



---

# Index

## A

- about this document xvi
- Access File Summary report 982
- access load reports
  - access time reports 892
  - creation time reports 893
  - modification time reports 893
  - overview 892
- access modes (MDisks) 130
- access privileges
  - granting 767
  - modifying 767
  - removing 768
  - viewing 767
- Access Time Summary report 981
- accessibility features
  - for Tivoli Storage Productivity Center 1213
- accessibility features for users with disabilities 1215
- accessing the server 638
- active management servers
  - about 643
  - reconnecting 651
  - taking over 651
- adding 622
  - CIM agents 243
  - copy sets to a FlashCopy session 725
  - copy sets to a Global Mirror session 729
  - copy sets to a Metro Global Mirror session 730
  - copy sets to a Metro Mirror session 728
  - copy sets to a Snapshot session 726
  - ESS/DS series logical paths 668
    - using a CSV file 669
  - host system connections 666
  - SNMP managers 652
  - storage connections 660
  - storage system 284
  - superuser groups 766
- adding element managers with CIMOM discovery 606
- administering users 763
- advanced options, Preferences menu 43
- Agent Manager certificates 128
- agents
  - changing the Windows service logon 243
  - checking the health 242
  - CIM 119
  - CIM agents 243
  - configuring for Storage Resource agent upgrade 254
  - deployment considerations for Storage Resource agents 120, 269
  - disabling automatic upgrade 266
  - enabling automatic upgrade 266

- agents (*continued*)
  - in-band Fabric agents 274, 276
  - migrating 252
  - migrating to Storage Resource agents 252
  - out-of-band agents 276
  - preparing Storage Resource agents for installation 120, 269
  - registering with a different server 268, 269
  - removing data 278
  - upgrading Storage Resource agent 254
  - viewing and managing 276
- alert
  - performance threshold 150
  - working with 383
- alert log 398, 401, 408
- alert notification 110, 1156
- alert overlay 399, 407, 408, 535
- alert tab 400, 408, 520, 536
- alerting
  - alert log 398, 401
  - computer alerts 383
  - database alerts 388
  - directory alerts 384
  - fabrics 390, 391, 393
  - file system alerts 385
  - instance alerts 387
  - notifications 110, 1156
  - other NAS alerts 386
  - overview 110, 1156
  - running scripts 507
  - table alerts 389
  - tables alerts 388
- alerts
  - database connection 393
  - endpoint device 393
  - fabric 390
  - performance thresholds 391
  - switch 391
  - triggering conditions 1158
  - viewing 398
- alias 115
- analysis report for Storage Optimizer 437, 438
- Analytics
  - Storage Optimizer 436
    - creating analysis report 437
    - creating optimization report 439
    - deleting an optimization report 442
    - deleting analysis jobs 442
    - retrieving optimization report job definition 443
    - viewing analysis report 438
    - viewing optimization report 441
- architecture
  - description 12
- archive job 144
- archive/backup actions 477

- archive/backup actions (*continued*)
  - creating job definitions 898
  - defining within constraints 477
  - overview 477
  - Reports 898
- arrays
  - performance reports 999
- asset reporting
  - agents 819
  - by cluster 808
  - by computer 811, 938
  - by DBMS type 938
  - by OS type 814
  - by storage subsystems 815
  - by Storwize V7000 Unified 821
  - information in asset reports 937
  - overview 807, 937
  - system-wide 818, 942
- Assign Host Ports wizard 573, 582
- assigning host ports 582
- audience xiii
- availability reporting
  - overview 831
- average volume utilization
  - storage virtualizers 1116
  - VDisk workload activity details reports 1123
  - VDisk workload activity reports 1110, 1113, 1120, 1126
- awareness, site 660, 722

## B

- backing up
  - considerations for 639
  - database 641
- backup
  - Tivoli Storage Productivity Center 181
  - Tivoli Storage Productivity Center for Replication 185
- backup job 144
- backup reporting
  - backup storage requirements reports 902
  - modified files not backed up reports 902
  - Most at Risk Files report 901
  - overview 901
- backup storage requirements reports
  - overview 902
- Basic HyperSwap
  - about 84, 682
  - session commands 708
- batch reports
  - creating 1021
  - creating for performance 1029
  - exporting and importing 1034
  - specifying output formats 1024
- Brocade 1143

- Brocade switches
  - discovery 117
  - entering user ID and password 333
- browsers supported 45

## C

- calculating read I/O capability
  - MDisk groups 1129
  - pools 1129
- capacity reporting 832
  - overview 948
- certificates
  - Agent Manager 128
  - importing SSL certificates for VMware 213
- change overlay 403
- changes in this release xix
- chargeback
  - overview 625
  - purpose 625
  - setting up invoices 627
- charts
  - types 787
- checking configuration 407
- CIM agents 119, 243
  - adding 244
  - collecting logs 249
  - discover automatically 245
  - removing 249
  - show managed devices 249
  - testing a connection 248
  - updating user identification 248
  - using Service Location Protocol 245
  - verify they are running 249
  - viewing information 246
- CIMOM 324
  - discovery 323
- Cisco 1143
- Cisco MDS 9000 118
- client ports 644, 653
- cluster (SAN Volume Controller)
  - access modes 130
  - overview 129
- collapsing the navigation tree 28
- collection
  - overview 77
- commands
  - repocopy 218
- comments, sending xvi
- communication-failure trap
  - descriptions 648
- components 12
- computer alerts 383
  - overview 383
  - triggering conditions 1158
- Computer Disk Space report 989
- computer quotas 452
- computers
  - discovery 330
- conditions
  - associating with groups 362
- configuration 422
  - analysis, using configuration 407
  - analyzing changes to 403, 405, 407
  - checking compliance to policies 407
  - history, using 403
- configuration change, SNMP trap descriptions 647
- configuration probe 422
- configuring
  - logical paths 668
  - Storage Resource agent for upgrade 254
  - storage systems 664
- Connection menu 41
- connections
  - about 656
  - adding
    - host systems 666
  - direct 657
  - Hardware Management Console (HMC) 658
  - modifying properties 663
  - removing 662
  - z/OS 658
- connections, displaying in topology viewer 547
- considerations, back up and recovery 639
- consistency group 676
  - time interval 676
- constraint 144
- constraint violation report 144
- constraint violations
  - affected hosts 973
  - affected volumes 973
  - creating reports 972
  - report 900, 979
  - viewing 973
- constraints
  - archive/backup actions
    - overview 477
  - creating 471
  - defining archive/backup actions 477
  - editing 474
  - overview 471
- content pane
  - overview 28
  - using 39
- controller cache performance
  - performance reports 1001
- controllers
  - performance reports 1002
- copy sets
  - about 671
  - adding 725, 726, 728, 729, 730
  - exporting 642, 734
  - importing 642, 734
  - removing 735
  - supported number of role pairs and volumes 671
- Create Virtual Disk wizard 574
- creating
  - constraint violation reports 972
  - file system extension policy 481
  - FlashCopy session 725
  - Global Mirror session 729
  - Metro Global Mirror session 730
  - Metro Mirror session 728
  - profile 354
  - Snapshot session 726
  - thin provisioning profiles 563
  - TPC server probes 341

- creating (*continued*)
  - VDisks 574
  - workload profiles 556

## D

- dashboard 31
- Data
  - system reports 981
- data agents 329
- Data agents
  - migrating 252
  - registering with a different server 269
  - uninstalling using GUI 266
- data collection 77
- data copying symbols 743
- data exposure 676
- Data Manager
  - overview 22
  - reports 804
- Data Manager for Chargeback
  - generating invoices 628
  - overview 24
  - planning 626
  - purpose 24
  - reports 961
  - viewing invoices 629
- Data Manager for Databases
  - reports 933
  - running scripts with alerts 507
  - system reports 993
  - what is it? 23
- data migration utility 736
- data path
  - monitoring and troubleshooting 538
- Data Path Explorer
  - overview 538
  - prerequisites 538
  - shortcut to launching 516
- Data Path Explorer view 545
- Data server
  - broadcasting location 194
  - configuring audit logging 194
  - displaying services 192
  - increasing memory 203
  - increasing memory on AIX 203
  - increasing memory on Linux 204
  - increasing memory on UNIX 204
  - increasing memory on Windows 204
  - shutting down services 195
  - starting a server service 195
- Data Server
  - changing DB2 password 208
  - checking status 202
- Data server service 191
- data sources
  - IBM Tivoli Storage Productivity Center servers 340
- database
  - backing up 641
  - restoring 641
- database alerts 388
  - description 388
- database connection alerts 393
- database repository
  - accessing views 217

- database repository (*continued*)
  - views overview 215
- database storage by computer
  - report 993, 994, 996, 997
- database-tablespace alerts
  - triggering conditions 1161
- database-tablespace quotas 459
- DB2 638
  - command line 182
  - db2cc command 182
  - monitoring 183
  - starting 637
  - stopping 637
- DB2 password
  - changing for Data server 208
  - changing for Device server 207
- default
  - groups 370
  - zone 115
- deleting
  - MDisks from an MDisk group 132
  - provision profiles 562
  - storage resource group sessions 434, 506
  - thin provisioning profiles 563, 564
  - VDisks 578
  - volumes 585
  - workload profiles 558
- Device server
  - changing DB2 password 207
  - changing host authentication password 208
  - changing WebSphere password 207
  - checking status 202
  - increasing memory 205
  - setting timeout values 205
- device triggering conditions 391
- directories reports
  - largest directories reports 899
  - monitored directory space reports 899
  - overview 899
- directory alerts 384
  - description 384
  - triggering conditions 1160
- directory groups
  - creating 373
  - overview 373
- disabilities, accessibility features for users
  - with 1215
- disabling
  - Metro Mirror heartbeat 734
  - Open HyperSwap 108, 706
  - OpenSwap 108, 706
- disaster recovery, practicing 739
- discovery 324
  - CIMOM 323
  - computers 330
  - fabrics 332
  - NetWare filers 330
  - storage resources 322
  - tasks 321
- Disk Defect report 984
- Disk Manager
  - overview 24, 25
  - reports 961
  - Storage Optimizer 436

- Disk Manager (*continued*)
  - creating analysis report 437
  - creating optimization report 439
  - deleting an optimization report 442
  - deleting analysis jobs 442
  - retrieving optimization report job definition 443
  - viewing analysis report 438
  - viewing optimization report 441
- Disk Space Summary report 982
- disk zone
  - definition 129
  - illustration 129
- Distributed Management Task Force (DTMF) 1208
- domain account
  - how to grant local administrative privileges 213
- domain port zoning 115
- Domain\_ID 115
- drop-down lists, limitations of to the sight-impaired 1215
- DS6000
  - port thresholds 1182
- DS8000
  - port thresholds 1182
  - Starting 602
  - using element managers to manage
    - adding element managers with config utility 600
    - modifying element managers 602
    - removing element managers 603
    - testing connections for element managers 603

**E**

- editing
  - file system extension policy 488
  - storage usage quotas 462
  - storage usage quotas for databases 464
- element managers
  - defining 566
  - launching 566
  - launching in topology viewer 548
  - working with DS8000 element managers 604
  - working with other element managers 604, 605
- email alert notification 110, 1156
- enabling, Metro Mirror heartbeat 734
- endpoint device alerts
  - creating 393
  - triggering conditions 1169
- enhancements to this release xix
- entities in topology viewer
  - arranging in a group 527
  - displaying information about 547
  - labels, specifying 548
  - locating 550
  - pinning 551
  - removing from database 552
  - selecting 553
- ESS/DS storage systems
  - adding logical paths 668

- ESS/DS storage systems (*continued*)
  - using a CSV file 669
  - removing logical paths 670
- expanding the navigation tree 28
- exporting and importing
  - batch reports 1034
- exporting copy sets 642, 734
- extents
  - capacity of cluster, effects on 132
  - definition 130
  - MDisk group, property of 130
  - overview 132
- external tools 622, 623, 624
  - overview 622

## F

- fabric 115
- Fabric agents
  - changing host authentication password 208
  - migrating 252
  - removing data 278
- fabric alerts
  - creating 390
  - triggering conditions 1167
- Fabric Manager
  - alerting 390, 391, 393
  - configure zones 117
  - create fabric groups 372
  - Fabric Manager user task summary 76
  - overview 26
  - probes 337
  - reports 973
  - system reports 1016
  - zone 115
  - zone set 115
- Fabric Performance Report - By Port 978
- Fabric zone configuration 1143
- fabric.mib file 1210
- fabrics 417, 422
  - alerting 390, 391, 393
  - create performance monitors 974
  - disabling 264
  - discovery 332
  - enabling 264
  - groups 372
  - probes 337
- failback 670
- failover 670
- FCID zoning 115
- fibre channel
  - methodologies of interconnects 1212
- file filters
  - wildcards 363
- File menu 40
- file system alerts
  - description 385
  - triggering conditions 1159
- file system extension
  - creating 481
- files reports
  - duplicate files reports 894
  - file size distribution reports 896
  - file summary reports 897
  - file types reports 897

- files reports (*continued*)
  - largest files reports 894
  - most obsolete files reports 895
  - orphan files reports 896
  - overview 894
- filesystem extension
  - editing 488
  - overview 480
- filesystem quotas 454
- filtering resources in topology viewer 548
- FlashCopy
  - about 85, 683
  - creating a session 725
  - session commands 708
- full disk encryption
  - icons, storage 521
  - support for 1145
- fully qualified host name
  - checking for on AIX 211
  - checking for on Linux 212
  - checking for on Solaris 212
  - checking for on Windows 211

## G

- Global Mirror
  - about 93, 691
  - creating a session 729
  - session commands 713
  - session details, viewing 760
- Global Mirror with Practice, session commands 715
- glossary 1223
- granting
  - access privileges 767
- groups 801
  - creating a storage subsystem 565
  - creating invoices for 368
  - Data Manager for Chargeback 368
  - default objects 370
  - deleting a storage subsystem 565
  - editing a storage subsystem 566
  - fabrics 372
  - monitoring 368
  - storage subsystem 565
- groups in topology viewer
  - creating custom 536
  - default 517
- groups, using 801
- GUI
  - content pane 28
  - menu bar 40
  - navigation tree 28
  - starting
    - from Tivoli Integrated Portal 54, 186
    - stand-alone 56, 188
  - tool bar 47
  - using 27

## H

- hardware zoning 115
- HBA 417, 422
- health, viewing summary of 742

- heartbeat, enabling and disabling 734
- Help menu 44
- historical trending 790
  - configuration history 403
  - exporting history data to .csv 793
  - history charts 790
- host authentication password
  - changing for Fabric agents 208
  - changing for the Device server 208
- host planner wizard 112
- host port 417, 422
- host ports
  - assigning to a VDisk 573
  - assigning to a volume 582
- host systems
  - adding a connection to 666
  - managing 666
  - removing 667
- host zone
  - definition 129
  - illustration 129
- hosts 417, 422
- hypervisor alerts
  - triggering conditions 1169

## I

- I/O group 129
- I/O groups
  - performance reports 1003
- I/O utilization
  - scenario for 137
- IBM 2005 switches 117
- IBM 2109 switches 117
- IBM Tivoli Enterprise Console event alert notification 110, 1156
- IBM Tivoli Storage Productivity Center
  - architecture 12
  - common user task summary 74
  - component overview 12
  - Data Manager user task summary 75
  - Disk Manager user task summary 75
  - menu bar 40
  - overview 11
  - tool bar 47
- IBM Tivoli Storage Productivity Center for Replication
  - server, accessing 638
  - starting 635
  - stopping 636
- IBM Tivoli Storage Productivity Center for Replication Three Site Business Continuity overview 51
- IBM Tivoli Storage Productivity Center for Replication Two Site Business Continuity overview 50
- IBM Tivoli Storage Productivity Center servers
  - overview 340
- icons
  - session status 743
  - storage systems status 655
- Identifying and investigating a suspended Metro Mirror 166
- identifying the version of IBM Tivoli Storage Productivity Center for Replication 639

- illustrations
  - disk and host zones 129
  - MDisk group 132
  - MDisk groups and VDIs 135
  - nodes (SAN Volume Controller) 129
- image VDisk 135
- importing copy sets 642, 734
- in-band discovery
  - information gathered 325
- installing
  - Storage Resource agent
    - considerations 120, 269
  - Storage Resource agents
    - remotely 120, 269
- instance alerts 387
  - overview 387
  - triggering conditions 1160
- instance quotas 457
- instance storage by computer report 998
- interface
  - content pane 39
  - navigation tree 39
  - navigation tree nodes 29
  - using 27
- Internet Engineering Task Force (IETF) 1210
- invoice 625, 626
  - samples 629
- invoices
  - generating 628
  - setting up 627
  - using the Online Execution window 628
  - viewing 629

## J

- jobs
  - managing 48
- journal volume, restoring data from 665

## L

- launch in context
  - available windows 1154
  - context task parameters 1153
  - parameters 1151
  - starting from the command line 63, 613
  - starting from Tivoli Integrated Portal 54, 186
  - starting with Java Web Start 61, 612
  - using outbound 68, 618
- Launch Planner 422
- LDAP user authentication
  - configuring in Tivoli Integrated Portal 226
- levels of detail in topology viewer 515
- libraries
  - probe 338
  - tape 371, 381, 383, 597, 598
  - tape groups 371
- licenses 21
- locating a tape cartridge 339
- locations
  - modifying for a site 735

- locations (*continued*)
  - modifying for a storage system 663
- logical paths
  - adding 668
    - using a CSV file 669
  - configuring 668
  - removing 670
  - viewing 668, 762
- logical SANs (LSANs), viewing topology 515
- logical storage subsystems
  - viewing volumes 762
- Login ID 110, 1156
- login notification alert notification 110, 1156
- LSANs (logical SANs), viewing topology 515

## M

- managed disk groups
  - performance reports 1005
  - workload activity reports 1103, 1107
- Management Information Base (MIB) files 1210
- management servers
  - about 643
  - reconnecting 651
  - setting as the standby server 649
  - setting up a remote server as the standby server 649
  - state change SNMP trap descriptions 648
  - taking over 651
  - viewing health summary 742
- managing
  - host systems 666
  - storage subsystems 565
  - storage systems 653
  - VDisks 573
  - volumes 582
- managing schedules, runs, and jobs 48
- McData 1143
- McDATA Intrepid 10000
  - removing data 278
- Mdisk Groups
  - VDisk workload activity reports 1110, 1113, 1126
- MDisk groups
  - adding an MDisk 573
  - creating 132
  - extents 130, 132
  - illustration 132, 135
  - MDisks, adding 132
  - MDisks, deleting 132
  - overview 132
  - performance reports 1005
  - states 132
  - workload activity reports 1103, 1107
- MDisks
  - adding a managed disk 574
  - adding to a MDisk group 573, 574
  - definition 129, 130
  - illustration 130
  - states 131

- memory
  - increasing allocation for Data server
    - on AIX 203
    - on Linux 204
    - on Windows 204
- memory allocation
  - Data server 203
  - increasing allocation for Device server 205
- menu bar 40
  - Connection menu 41
  - File menu 40
  - Help menu 44
  - Preferences menu 41
  - View menu 41
  - Windows menu 43
- messages, viewing 763
- metricsperformance reports 1186
- Metro Global Mirror
  - about 98, 696
  - creating a session 730
  - session commands 716
- Metro Global Mirror scenarios
  - disaster recovery (unplanned outages) 171
  - planned outages 168
  - practice 175
- Metro Global Mirror with Practice, session commands 719
- Metro Mirror
  - about 87, 685
  - creating a session 728
  - session commands 710
- Metro Mirror heartbeat, enabling and disabling 734
- Metro Mirror with Practice, session commands 712
- migrating
  - Data agent 252
- minimap in Configuration History view 407
- minimap in topology viewer 516
- modified files not backed up reports 902
- modifying 623
  - access privileges 767
  - provisioning profiles 562
  - site location 735
  - storage connection properties 663
  - storage system location 663
  - thin provisioning profiles 563, 565
  - workload profiles 558
- monitor
  - creating a storage-subsystem group 565
  - deleting a storage-subsystem group 565
  - editing a storage-subsystem group 566
- monitoring 1176
  - groups 368
  - performance with Data Path Explorer 538
  - tape libraries 381, 598
- Monitoring Agent
  - displaying reports 126, 1130

- Most at Risk Files report 901, 986
- Most Obsolete Files report 987
- multipath 417, 422

## N

- N\_Port address identifier 115
- N\_Port\_Name 115
- naming conventions
  - VDisks 574
- NAS alerts
  - triggering conditions 1169
- navigation tree
  - collapsing 28
  - expanding 28
  - overview 28
  - using 39
- NetWare
  - scanning NetWare volumes 351
- NetWare filer 329
- NetWare filers
  - discovery 330
- Network Appliance OS user quotas
  - viewing 470
- Network Appliance quotas
  - creating 467
  - viewing directory (QTree) quotas and violations 470
- Network Appliance user quotas
  - viewing 469
- network quotas 447, 449
- Node\_Name 115
- nodes (SAN Volume Controller)
  - illustration 129
  - overview 129
- non standard zones
  - fabric assist 118
  - protocol 118
  - quick loop 118

## O

- objects
  - default 370
- Oldest Orphaned Files report 988
- Online Execution window 628
- Open HyperSwap
  - disabling 108, 706
- OpenSwap
  - disabling 108, 706
- operating system user authentication
  - configuring in Tivoli Integrated Portal 238
- optimization report for Storage Optimizer 439, 441
- orphan zones 117
- OS Disk Capacity By Computer
  - sample invoice 630
- OS Storage Usage By User
  - sample invoice 629, 630
- OS user groups reports 899
- other NAS alerts 386
  - overview 386
- out-of-band discovery
  - information gathered 325

- overview
  - IBM Tivoli Storage Productivity Center for Replication Three Site Business Continuity 51
  - IBM Tivoli Storage Productivity Center for Replication Two Site Business Continuity 50

## P

- pager alert notification 110, 1156
- passwords
  - changing 206
  - changing DB2 password for Data server 208
  - changing DB2 password for Device server 207
  - changing host authentication password for Fabric agents 208
  - changing host authentication password for the Device server 208
  - changing using password tool 209
  - changing WebSphere password for the Device server 207
- path planning 417, 422
- paths
  - adding 668
    - using a CSV file 669
  - removing 670
- PDF file, saving as 785
- performance 1175
  - monitoring with Data Path Explorer 538
- performance monitor 417, 422
- performance monitor, creating 974
- performance overlay 407
- performance reports
  - array performance 999
  - controller cache performance 1001
  - controller performance 1002
  - creating 968
  - describing column metrics
    - back-end 1186
    - front-end and fabric 1186
    - volume 1186
  - I/O groups 1003
  - managed disk groups 1005
  - port performance 1007
  - predefined reports 999
  - SAN Volume Controller node cache 1006
  - storage subsystem performance 1009
  - top n active volumes cache hit performance 1010
  - top volumes data rate 1012
  - top volumes disk performance 1012
  - top volumes I/O rate 1014
  - top volumes response performance 1015
- performance thresholds 391, 1176
- performance triggering conditions 391
- performing
  - takeover 651
- physical disks
  - definition 130
  - illustration 130

- pings
  - creating 364
  - pings 364
- pinning entities in topology viewer 551
- planned outage of remote (H3) and return 171
- planned outage of the local site (H1) while in normal configuration, with a production move to the remote site (H3) 170
- planned outage of the local site (H1), with a production move to the intermediate site (H2) and an optional return 169
- planned outage on site 2 (H2) and return three sites to the original configuration 169
- planner
  - provisioning 111
  - security 114
  - storage 114
- planning
  - for Data Path Explorer view 545
- policy management
  - constraints 471
  - creating computer quotas 452
  - creating constraints 471
  - creating database-tablespace quotas 459
  - creating filesystem quotas 454
  - creating instance quotas 457
  - creating Network Appliance quotas 467
  - creating network quotas 447
  - creating network quotas for databases 449
  - editing constraints 474
  - editing storage a usage quota 462
  - editing storage a usage quota for databases 464
  - enforcing storage policies 444
  - filesystem extension 480
  - quotas 446, 466
  - scheduled actions 493
  - viewing imported Network Appliance OS user quotas 470
  - viewing imported Network Appliance user quotas 469
  - viewing Network Appliance directory (QTree) quotas and violations 470
- Port Connections report 1016
- port name zoning 115
- port performance
  - performance reports 1007
- port zoning 115
- PortId zoning 115
- ports 644, 653
- practice at H3 for Metro Global Mirror 177
- practice recover to H1 with production continuing on H2 176
- practice recover to H2 with production continuing on H1 in a normal configuration 175
- practice volumes 675, 739
- practicing disaster recovery 739
- predefined performance reports 999

- Preferences menu 41
- Preserve Mirror option, recommendations for 723
- primary server, reinstalling during active session 650
- primordial 520
- printing 785
- probe 144, 417
- PROBE\_ME file 269
- probes
  - creating 336
  - probes 336
  - fabric 337
  - monitoring probes 335
  - overview 335
- profile management 558
- profiles
  - associating conditions with groups 362
  - creating 354
  - using wildcards 363
- properties
  - session 748
  - viewing session 760
- properties, user-defined 567
- protected volumes, about 659
- protecting volumes 664
- provisioning planner 111
- provisioning profiles
  - deleting 562
  - modifying 562
  - viewing 562
  - working with 558

## Q

- QLogic 1143
- quota violations reporting
  - all quotas reports 901
  - computer quotas reports 901
  - file system quotas reports 901
  - network-wide quotas reports 901
- quotas
  - creating computer quotas 452
  - creating database-tablespace quotas 459
  - creating filesystem quotas 454
  - creating instance quotas 457
  - creating Network Appliance quotas 467
  - creating network quotas 447
  - creating network quotas for databases 449
  - editing a storage usage quota 462
  - editing a storage usage quota for databases 464
  - overview 446, 466
  - viewing imported Network Appliance OS user quotas 470
  - viewing imported Network Appliance user quotas 469
  - viewing Network Appliance directory (QTree) quotas and violations 470



## R

- RAID 417
  - volumes, setting for 583
- read I/O capability
  - formula 1129
- reader feedback, sending xvi
- recommendations for using the Preserve Mirror option 723
- reconnecting the active and standby management servers 651
- recovering, considerations for 639
- recovery
  - Tivoli Storage Productivity Center for Replication 185
- refreshing, storage system configuration 664
- registering the Data agent with a different server 269
- removing 624
  - copy sets 735
  - host systems 667
  - logical paths 670
  - sessions 735
  - storage connection 662
  - storage system 292
  - storage systems 662
  - user access 768
- replication planning 417
- repocopy command 218
  - exporting data 218
  - importing data 219
- repocopy tool
  - exporting data 218
- report
  - constraint violation 972
  - storage subsystem 966
  - storage subsystem constraint violation 972
  - storage subsystem performance 968
- reporting 801
  - archive/backup 898
  - asset reports 807, 937
  - availability 831
  - backup 901
  - by clauses 806, 936
  - capacity 832
  - capacity reports 948
  - categories 769, 771
  - considerations 781
  - dashboard 31
  - drilling down 789
  - generating reports 793
  - generating rollup reports 1036
  - historical trending 790
  - history charts 790
  - limitations 781
  - overview 769, 771
  - printing 785
  - reporting
    - categories 769
    - selecting an object 797
    - sorting columns 785
    - tabs, navigating 789
    - types 769
    - usage 892
    - usage reports 954
    - usage violation 899, 959

- reporting (*continued*)
  - what can I find out? 777
- reports 980
  - Data Manager 804
  - Data Manager for Chargeback 961
  - Data Manager for Databases 933
  - Disk Manager 961
  - Fabric Manager 973
  - system 981, 993
  - Tape Manager 980
  - VDisk workload activity details reports 1123
  - VDisk workload activity reports 1110, 1113, 1120, 1126
  - VMware capacity 932
- Repository Copy tool 218
  - importing data 219
- restoring data from a journal volume 665
- restoring, database 641
- role pairs 675
- role pairs, status messages 759
- roles, user
  - description of 71
- runs
  - managing 48

## S

- SAN
  - Device server 324
- SAN (storage area network), viewing topology 515
- SAN Assets (ALL) report 1017
- SAN Assets (Connected Devices) report 1017
- SAN Assets (Switches) report 1017
- SAN Planner 417, 422
- SAN Planner Wizard
  - Editing the plan 433
  - Executing the plan 433
  - Reviewing user selections 432
  - Selecting replication sessions 429
  - Selecting the planning task 423
  - Selecting the provision plan details 423
  - Specifying a secondary location 431
  - Specifying a tertiary location 432
  - Specifying replication session properties 430
  - Specifying selected volume assignments for an SRG 425
  - Specifying selected volume assignments for input resources 425
  - Specifying the advanced capacity plan 427
  - Specifying the capacity plan 426
  - Specifying the path plan 428
  - Specifying the plan content 423
  - Specifying the zone plan 429
- SAN Volume Controller
  - extents, definition 130
  - I/O group 129
  - illustration 129
  - MDisk group 132
- SAN Volume Controller (*continued*)
  - MDisks
    - access modes 130
    - definition 129, 130
    - illustration 130
  - nodes 129
  - overview 129
  - physical disks
    - definition 130
    - illustration 130
  - publications xiii
  - VDisks 129
- SAN Volume Controller nodes
  - cache performance reports 1006
- scans
  - creating 344
    - scans 344
  - for archive 144
  - monitoring
    - scans 343
  - NetWare volumes 351
  - overview 343
- scenario
  - Identifying and investigating a suspended Metro Mirror 166
- scenarios
  - adding storage to a server and setting up a zone 589
  - creating volumes and assigning them to a host 583
  - identifying performance hot spots using Storage Optimizer 161
  - locating a tape cartridge 339
  - monitoring and reporting on I/O utilization 137
  - overview of scenarios for practicing at H3 for Metro Global Mirror 177
    - practice scenario 1 178
    - practice scenario 2 178
    - practice scenario 3 179
  - planned outage of remote (H3) and return 171
  - planned outage of the local site (H1) with a production move to the remote site (H3) 170
  - planned outage of the local site (H1), with a production move to the intermediate site (H2) and an optional return 169
  - planned outage on site 2 (H2) and return three sites to the original configuration 169
  - practice at H3 for Metro Global Mirror 177
  - practice recover to H1 with production continuing on H2 176
  - practice recover to H2 with production continuing on H1 in a normal configuration 175
  - retiring systems using Storage Optimizer 163
  - unplanned outage of local (H1) with production move to remote (H3) 174
  - Unplanned outage of local (H1), production move to intermediate (H2) 171

- scenarios (*continued*)
  - unplanned outage of remote (H3) 173
  - unplanned outage of Site 2 (H2) 172
- scenarios for Metro Global Mirror 168, 171, 175
- scenarios overview for practicing at H3
  - for Metro Global Mirror 177
  - practice scenario 1 178
  - practice scenario 2 178
  - practice scenario 3 179
- scheduled actions
  - create script window 494, 508
  - creating 494, 508
  - edit script window 496
  - editing 496
  - overview 493
- schedules
  - managing 48
- screen shots
  - Create Virtual Disk wizard 574
  - Storage Subsystem window 566
  - Storage-Subsystem Details page 567
  - Virtual Disks page 578
  - Virtual-Disk Details page 578
  - Volumes page 582
- scripts
  - resolving script names 512
  - running 510
  - running scripts with alerts 507
  - using with jobs 507
- SDD 417
- Secure Socket Layer (SSL)
  - disabling 237
  - enabling 234
- security 763
  - considerations 127
- security planner 114
- segments with wasted space report 995
- sending comments xvi
- sequential VDisk 135
- service
  - Data server
    - starting 195
  - Data Server 191
- Service Location Protocol
  - configuration recommendations 324
  - using to discover CIM agents 245
- Service Location Protocol (SLP)
  - overview 1210
- Service tool
  - for agents 222
  - for servers 222
- services
  - Data server
    - broadcasting location 194
    - configuring audit logging 194
    - configuring tracing 193
    - configuring tracing services 193
    - displaying 192
    - displaying service log 193
    - shutting down 195
  - Device server 196
    - changing the password 197
    - displaying information 198
    - viewing log 198
  - starting 200
- services (*continued*)
  - stopping 198
  - viewing and managing 190
- session commands
  - Basic HyperSwap 708
  - FlashCopy 708
  - Global Mirror 713
  - Global Mirror with Practice 715
  - Metro Global Mirror 716
  - Metro Global Mirror with Practice 719
  - Metro Mirror 710
  - Metro Mirror with Practice 712
  - Snapshot 709
- session state change trap
  - descriptions 646
- session types
  - Basic HyperSwap 84, 682
  - FlashCopy 85, 683
  - Global Mirror 93, 691
  - Metro Global Mirror 98, 696
  - Metro Mirror 87, 685
  - Snapshot 86, 684
- sessions
  - about 670
  - FlashCopy 725
  - Global Mirror 729
  - images 743
  - Metro Global Mirror 730
  - Metro Mirror 728
  - properties 748
  - removing 735
  - Snapshot 726
  - states 745
  - status icons 743
  - viewing 743
  - viewing details 760
    - Global Mirror 760
  - viewing health summary 742
  - viewing properties 760
- setting
  - remote standby management server 649
  - SNMP 652
  - standby management server 649
  - volume protection 664
- setting units of measurement for reports 43
- setting up invoices 627
- shortcuts in minimap 516
- sight-impaired 1215
- single sign-on
  - overview 58
  - user authentication method used with 58
- site awareness 660, 722
- site, modifying location of 735
- SMI-S support 324
- Snapshot
  - about 86, 684
  - creating a session 726
  - session commands 709
- snapshot selection panel 403
- SNMP
  - adding managers 652
  - MIBs 1210
  - overview 1210
- SNMP (*continued*)
  - setting up 652
  - trap alert notification 110, 1156
  - traps 1210
  - viewing alerts 742
- SNMP trap descriptions
  - communication-failure 648
  - configuration change 647
  - management servers state change 648
  - session state change 646
  - suspending-event notification 647
- SNMP traps
  - in-band 324
- software zoning 115
- solid-state drives
  - icons, storage 521
  - support for 1145
- SSL certificates
  - how to import for VMware 213
- standby management servers
  - about 643
  - ports 644, 653
  - setting 649
  - setting a remote server 649
- starting 624
  - DB2 637
  - DS8000 602
  - from Tivoli Integrated Portal
    - Tivoli Storage Productivity Center 54, 186
  - IBM Tivoli Storage Productivity Center for Replication 635
  - stand-alone
    - Tivoli Storage Productivity Center 56, 188
- starting services 200
  - Tivoli Integrated Portal server 202
- states, session 745
- status messages, role pairs 759
- status, checking in topology viewer 527
- status, session icons 743
- stopping
  - DB2 637
  - IBM Tivoli Storage Productivity Center for Replication 636
- stopping services 198
- Storage Access Times report 983
- storage area networks (SANs), viewing topology 515
- Storage Availability report 989
- storage capacity 961
- storage connections
  - removing 662
  - viewing details 762
- storage controller 422
- Storage Management Initiative - Specification (SMI-S) 1209
- Storage Modification Times report 985
- storage network, viewing 515
- Storage Networking Industry Association (SNIA) 1209
- Storage Optimizer
  - analysis report 435
  - creating analysis report 437
  - creating optimization report 439
  - deleting an optimization report 442

- Storage Optimizer (*continued*)
  - deleting analysis jobs 442
  - launching from topology viewer 550
  - optimization report 435
  - overview 435
  - retrieving optimization report job
    - definition 443
  - scenario for identifying hot spots 161
  - scenario for retiring systems 163
  - starting from topology viewer 550
  - supported storage systems 435
  - using 436
  - viewing analysis report 438
  - viewing optimization report 441
  - zoning recommendations provided
    - by 443
- storage planner 114
- storage pool 422
- storage pools 417
- storage port 417, 422
- Storage Resource agents
  - changing authentication 259
  - configuring for upgrade 254
  - deployment considerations 120, 269
  - registering with a different
    - server 268
  - uninstalling using GUI 266
  - upgrading 254
- storage resource group
  - VDisk workload activity details
    - reports 1123
  - VDisk workload activity reports 1120
- storage resource group sessions
  - deleting 434, 506
  - groups
    - storage resource groups 434, 506
    - viewing 434, 506
    - working with 434, 506
- storage resource groups 417
  - groups
    - storage resource groups 503
    - overview 503
- storage subsystem
  - creating a group 565
  - data collection 322
  - deleting a group 565
  - editing a group 566
  - group 565
  - performance reports 1009
- storage subsystem alerts
  - triggering conditions 1163
- storage subsystem constraint violation
  - reporting
    - generate report 972
- storage subsystem reporting 961
  - computer views 961
  - generate report 966
  - storage subsystem views 961
  - Volume to HBA assignment 961
- Storage Subsystem window, screen
  - shot 566
- storage subsystems
  - managing 565
  - overview 284
  - removing 567
  - setting user-defined properties 567
  - viewing information 568
- storage system
  - adding 284
  - removing 292
  - testing connection 292
- storage systems
  - about 655
  - adding a connection to 660
  - icons 655
  - managing 653
  - modifying connection properties 663
  - modifying location of 663
  - ports 644, 653
  - refreshing the configuration 664
  - removing 662
  - removing a connection to 662
  - supported 655
  - viewing details 761
  - viewing health summary 742
  - viewing volumes 762
- storage tier reports 1097
- storage tiering
  - reports 1103, 1107
- Storage Usage By Database User
  - sample invoice 631
- storage virtualizers
  - VDisk workload activity reports 1116
- storage-subsystem controllers
  - performance reports 1002
- Storage-Subsystem Details page, screen
  - shot 567
- storage-subsystem performance
  - reports 968
- striped VDisk 135
- summary of changes in this release xix
- superuser group
  - adding administrator role to 766
- suspending-event notification trap
  - descriptions 647
- switch 417, 422
- switch alerts
  - creating 391
  - triggering conditions 1167
- switch performance
  - gathering data 335
- Switch Performance Management 1143
- switch performance monitor
  - creating 375
  - editing 376
- Switch Performance report 977, 1018
- switch performance report, define and
  - generate 974
- Switch Port Errors report 1018
- switch zone 417, 422
- symbols
  - data copying 743
  - session status 743
  - volume roles 743
- symmetric virtualization 129
- SYSMANT authority
  - checking and setting 183
- system maintenance authority
  - checking and setting 183
- system reports 981, 993
  - Access File Summary 982
  - Access Time Summary 981
  - Computer Disk Space 989
- system reports (*continued*)
  - database storage by computer 993, 994, 996, 997
  - Disk Defects 984
  - Disk Space Summary 982
  - Fabric Manager 1016
  - instance storage by computer 998
  - Most at Risk Files 986
  - Most Obsolete Files 987
  - node cache performance 1006
  - Oldest Orphaned Files 988
  - Port Connections 1016
  - SAN Assets (ALL) 1017
  - SAN Assets (Connected
    - Devices) 1017
  - SAN Assets (Switches) 1017
  - segments with wasted space 995
  - Storage Access Times 983
  - Storage Availability 989
  - Storage Modification Times 985
  - Switch Performance 977, 1018
  - Switch Port Errors 1018
  - Top Switch Ports Data Rate
    - Performance 1019
  - Top Switch Ports Packet Rate
    - Performance 1020
  - total database freespace 994, 996, 997
  - total device freespace 997
  - total DMS container freespace 998
  - Total Freespace 991
  - user database quota violations 993
  - user database space usage 993, 994, 996, 998, 999
  - User Quota Violations 988
  - User Space Usage 992
  - using 980
  - Wasted Space 992

## T

- T11 FC-MI 1212
- table alerts 389
  - description 389
  - triggering conditions 1162
- tablespace alerts 388
  - description 388
- tabs
  - alert 520
- tabular view
  - tape libraries 554, 555
- taking over management from the active
  - server 651
- tape
  - topology 554
  - topology viewer
    - tape libraries 554
- tape cartridge, locating 339
- tape libraries 980
  - alerts 402
  - discovering 597
  - groups 371
  - managing 596
  - monitoring 338, 381, 598
  - probes 338
  - tabular view 554, 555
  - topology view 554
  - viewing details 383, 597

- Tape Manager
    - monitoring 338, 381, 598
    - overview 596
    - reports 980
  - temporary zone sets 115
  - testing connection
    - storage system 292
  - thin provisioning profiles
    - creating 563
    - deleting 563, 564
    - modifying 563, 565
    - viewing 565
    - working with 563
  - threshold
    - alert 150
    - array 1177
    - array threshold 1177
    - controller 1181
    - controller thresholds 1181
  - thresholds 1176
    - port 1182
    - port thresholds 1182
  - time zone 417, 422
  - timeout values
    - setting for Device server 205
  - Tivoli Common Reporting
    - reports 1097
  - Tivoli Integrated Portal
    - overview 53
    - starting 54, 186
    - starting server 202
  - Tivoli Storage Manager archive/backup
    - actions, see archive/backup
    - actions 477
  - Tivoli Storage Productivity Center
    - managing schedules, runs, and jobs 48
  - Tivoli Storage Productivity Center for Replication
    - accessing 185
    - backing up 185
    - recovering 185
    - server, accessing 185
  - Tivoli Storage Productivity Center servers
    - overview 278
  - tivoliSRM.mib file 1210
  - tool bar 47
  - tools
    - repository copy 218
  - top n active volumes cache hit
    - performance
      - performance reports 1010
  - Top Switch Ports Data Rate Performance
    - report 1019
  - Top Switch Ports Packet Rate
    - Performance report 1020
  - top volumes data rate
    - performance reports 1012
  - top volumes disk performance
    - performance reports 1012
  - top volumes I/O rate
    - performance reports 1014
  - top volumes response performance
    - performance reports 1015
  - topology 417, 422
  - topology settings, shortcut to 516
  - topology viewer
    - arranging entities within a group 527
    - connections, displaying 547
    - Data Path Explorer 538
    - element managers, launching 548
    - entities, removing from database 552
    - entities, selecting 553
    - entity information, displaying 547
    - filtering resources 548
    - groups, creating custom 536
    - groups, default 517
    - icons, storage 521
    - labels, specifying 548
    - launch planner 550
    - launch Storage Optimizer 550
    - levels of detail 515
    - locating entities 550
    - manage storage subsystems 551
    - manage zone control 551
    - minimap 516
    - missing entities 553
    - monitoring and troubleshooting a data path 538
    - pinning entities to monitor 551
    - reports 547
    - show alert logs 553
    - show monitoring status 554
    - start Storage Optimizer 550
    - status, checking 527
    - table view 520
    - tape libraries 554
    - user-defined properties (UDPs) 537
    - viewing a single data path 538
    - zones and zone sets 546
  - total database freespace report 994, 996, 997
  - total device freespace report 997
  - total DMS container freespace 998
  - Total Freespace report 991
  - Total Size By Database -Tablespace
    - sample invoice 632
  - TotalStorage Enterprise Storage Server
    - port thresholds 1182
  - TPC server probes
    - creating 341
  - trademarks 1220
  - translations
    - browser locale requirement xv
  - traps for SNMP 1210
  - trending 790
    - history charts 790
  - triggering conditions
    - alerts 1158
    - computer alerts 1158
    - database-tablespace alerts 1161
    - device 391
    - directory alerts 1160
    - endpoint device alerts 1169
    - fabric alerts 1167
    - file system alerts 1159
    - hypervisor alerts 1169
    - instance alerts 1160
    - NAS alerts 1169
    - performance 391
    - storage subsystem alerts 1163
    - switch alerts 1167
    - table alerts 1162
  - troubleshooting
    - too many host ports (SAN Volume Controller) 573
    - volumes not displayed 586
- ## U
- UDPs (user-defined properties) in topology viewer 537
  - uninstalling
    - Data agent 266
    - Storage Resource agent 266
  - units of measurement, setting for reports 43
  - UNIX Event alert notification 110, 1156
  - unplanned outage of local (H1) with production move to remote (H3) 174
  - Unplanned outage of local (H1), production move to intermediate (H2) 171
  - unplanned outage of remote (H3) 173
  - unplanned outage of Site 2 (H2) 172
  - upgrading agents 266
  - usage reporting
    - access load reports 892
    - directories reports 899
    - files reports 894
    - OS user groups reports 899
    - overview 892, 954
    - users reports 899
  - usage violation reporting
    - constraint violations report 900
    - overview 899, 900, 959
    - quota violations reports 900
  - Usage Violation Reporting
    - information in Usage Violation reports 960
    - Quota Violations 960
  - user administration 763
    - adding superuser groups 766
    - granting access privileges 767
    - modifying access privileges 767
    - removing user access 768
    - viewing access privileges 767
  - user authentication
    - configuring in Tivoli Integrated Portal
      - LDAP 226
      - operating system 238
    - Secure Socket Layer (SSL)
      - disabling 237
      - enabling 234
  - user database quota violations
    - report 993
  - user database space usage report 993, 994, 996, 998, 999
  - user group name
    - adding to Tivoli Storage Productivity Center for Replication 184
  - user interface
    - content pane 39
    - menu bar 40
    - navigation tree 39
    - navigation tree nodes 29
    - starting
      - using Tivoli Integrated Portal 54, 186, 202
    - tool bar 47

- user interface (*continued*)
  - using 27
- User Quota Violations report 988
- user roles
  - description of 71
- User Space Usage report 992
- user-defined properties
  - setting 567
- user-defined properties (UDPs) in topology viewer 537
- users
  - about 763
  - groups 763
  - roles 764
- users reports 899
- using practice volumes 739

**V**

- VDisks
  - assigning host ports 573
  - creating 574
  - definition 129
  - deleting 135, 578
  - illustration 135
  - managing 573
  - naming conventions 574
  - overview 135
  - states 135
  - types
    - sequential 135
    - striped 135
  - viewing information 578
- ver command 639
- version of IBM Tivoli Storage Productivity Center for Replication 639
- View menu 41
- viewing
  - access privileges 767
  - alerts 398
  - configuration changes in graphical view 405
  - configuration changes in table view 407
  - constraint violations 973
  - error messages 763
  - health summary 742
  - logical paths 668, 762
  - provisioning profiles 562
  - session details 760
    - Global Mirror 760
  - session properties 760
  - sessions 743
  - SNMP alerts 742
  - storage capacity 961
  - storage connection details 762
  - storage resource group sessions 434, 506
  - storage system details 761
  - storage-subsystem information 568
  - thin provisioning profiles 565
  - VDisk information 578
  - volume details 762
  - volume information 586
  - volumes affected by constraint violations 973

- viewing (*continued*)
  - workload profiles 558
- viewing invoices 629
- viewing your topology
  - arranging entities within a group 527
  - connections, displaying 547
  - element managers, launching 548
  - entities, removing from database 552
  - entities, selecting 553
  - entity information, displaying 547
  - filtering resources 548
  - groups, creating custom 536
  - groups, default 517
  - icons, storage 521
  - labels, specifying 548
  - levels of detail 515
  - locating entities 550
  - minimap 516
  - pinning entities to monitor 551
  - reports 547
  - status, checking 527
  - table view 520
  - user-defined properties (UDPs) 537
  - zones and zone sets 546
- views
  - overview 215
  - retrieving information 215, 217
- Virtual Disks page, screen shot 578
- virtual SANs (VSANs), viewing topology 515
- Virtual-Disk Details page, screen shot 578
- VMware
  - capacity reports 932
- VMware data source
  - viewing and managing 292
- VMware support
  - how to import SSL certificates 213
- volume creation 422
- volume planning 417, 422
- volume protection
  - about 659
  - setting 664
- volume roles
  - about 674
  - symbols 743
- volumes 583
  - assigning host ports 582
  - creating 583
  - creating volumes and assigning them to a host 583
  - deleting 585
  - managing 582
  - practice 675, 739
  - RAID levels, selecting 583
  - role pairs 675
  - roles 674
  - viewing details 762
  - viewing information 586
- Volumes Details page 586
- Volumes page, screen shot 582
- VPA 417, 422
- VSANs 118
- VSANs (virtual SANs), viewing topology 515

**W**

- Wasted Space report 992
- Web Based Enterprise Management (WBEM) 1208
- WebSphere password
  - changing for the Device server 207
- what's new in this release xix
- windows
  - create script 494, 508
  - edit script 496
  - instance alerts 387
  - Online Execution 628
  - Storage Subsystem window 566
  - Storage-Subsystem Details page 567
  - table alerts 389
  - Virtual-Disk Details page 578
  - Volumes page 582
- Windows log alert notification 110, 1156
- Windows menu 43
- wizard
  - host planner 112
- wizards
  - Assign Host Ports wizard 573, 582
  - Create Virtual Disk wizard 574
- working with
  - provisioning profiles 558
  - storage resource group sessions 434, 506
  - thin provisioning profiles 563
- workload profile 417, 422
- workload profiles
  - creating 556
  - deleting 558
  - modifying 558
  - viewing 558
- WWN zoning 115
- WWPN 417

**X**

- XIV Storage System
  - port thresholds 1182

**Z**

- zone 115, 417, 422
  - change properties 590
  - concepts 115
  - configure 117
  - creating 588
  - default 115
  - delete 591
  - duplicate 591
  - hardware 115
  - membership 115
  - non standard 118
  - orphan 117
  - port 115
  - software 115
  - supported switches 118
  - working with 588
- zone alias
  - change properties 594
  - creating 594
  - delete 596
  - duplicate 595

- zone alias (*continued*)
  - working with 594
- zone planning 417, 422
- zone set 115, 417, 422, 592
  - activate 593
  - change definition 592
  - creating 592
  - data source 593
  - deactivate 593
  - delete while contains zones 117
  - duplicate 593
  - inactive 593
  - temporary 115
- zones 115
- zones and zone sets in topology
  - viewer 546





Product Number: 5608-WB1, 5608-WC0, 5608-WC3, 5608-WC4, 5725-A01,  
5725-E18  
5608-WB2, 5608-WB3

Printed in USA

SC27-2338-08

