IBM Spectrum Control 5.4.0, August 2020

User's Guide





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

Contents

About this guide	
Who should read this guide	i>
Publications	i)
Accessing publications online	i)
IBM Redbooks)
Translation)
Providing feedback about publications	
IBM Spectrum Control technical community	
Contacting IBM Software Support	
Reporting a problem	
Conventions used in this guide	
What's new	xii
Change history	
Discontinued features	
Sponsor user program	
Beta program	
Deta program	
Chapter 1. Getting started	
Managing resources	
Configuring and administering.	
Navigation	
Chapter 2. Social media	
Videos	
Blogs	
51080	
Chapter 3. Product licenses	25
Definitions of licensing metrics	
Names of equivalent licenses in previous releases	
Actions that are available based on role	
Licensing	
How licensing values are calculated	
Configuring license values for storage systems	
Licensing examples	
Chamter 4. Draduct eversions	4.0
Chapter 4. Product overview	41
Interfaces for IBM Spectrum Control.	
Starting IBM Spectrum Control	
Navigation	
IBM Spectrum Storage Suite	
IBM Storage Insights	
Replication products	
Product updates and security fixes	
Releases and downloads	
Security	
Subscribing to IBM announcements	
End of support	
Collaborating with the team	63

Reporting a problem	63
Key concepts	
Data collection	
Performance monitoring and troubleshooting	
Applications, departments, and general groups	
Storage optimization	
Cloud configuration	
Provisioning storage	
Storage reclamation	
Alerts and alert policies	
Reporting	
Rollup servers	
Units of measurement for storage data	
Role-based authorization	
Fabrics and zones	
AgentsREST APIs	
Chantay E. Tutaviala	07
Chapter 5. Tutorials Optimizing the performance of storage virtualizers	
Tutorial: Analyzing and re-tiering volumes in pools on tier 1	
Tutorial: Collocating volumes	
Monitoring capacity usage at different levels of a business hierarchy	
Tutorial: Comparing storage usage in each department	
Using applications and subcomponents to monitor capacity and space usage	
Tutorial: Viewing storage capacity and usage trends	
Tutorial: Viewing NPIV connections between server ports and switch ports in a fabric	
Tutorial: Exporting performance data for a SAN Volume Controller system	
Tutorial: Comparing the performance of storage systems	
Tutorial: Reviewing and updating your agentless servers	
Tutorial: Troubleshooting performance	
Tutorial: Monitoring IBM Spectrum Scale performance	
Tutorial: Viewing the aggregated workload for an application	
Tutorial: Identifying the source of slow drain problems caused by depletion of buffer credits	
Chapter 6. Managing resources	111
Resources that you can monitor.	
Adding resources	
Required user roles for monitoring resources	
Adding storage systems	
Adding fabrics and switches	
Adding hypervisors	
Adding servers	134
Adding rollup servers	136
Restrictions for rollup servers	137
Removing resources	137
Collecting data	138
Collecting asset and status data	139
Collecting performance data	
Collecting information about shares on storage systems	
Collecting information about the sizes of snapshots in IBM Spectrum Scale	
Alerting	
How alerts work	
Viewing and administering alerts	
Viewing and administering alert definitions	
Alert policies	
Defining notification settings for alerts	157

Alert severities	158
Defining alerts	159
Configuring alert notifications	186
Triggering conditions for alerts	188
Alert notifications and actions	226
Viewing information about resources	231
How information is organized	231
Dashboard view	
Viewing information about top-level resources	236
Viewing information about internal, object, and related resources	
Overview charts	
Viewing information about enclosures	
Monitoring the status and condition of resources	
How the condition of a resource is determined	
Viewing the overall condition of resources	
Viewing the condition of specific types of resources	
Viewing the status of resources	
Acknowledging the condition and status of resources	
Monitoring vaults in IBM Cloud Object Storage	
Monitoring the performance of resources	
Viewing performance information	
Key performance indicators	
Viewing performance alerts	
Exporting performance data for storage systems and fabrics	
Performance metrics	
Monitoring the capacity of resources	
Setting capacity limits	
Viewing capacity information	
Viewing the capacity of external storage	
Viewing capacity alerts and violations	
Capacity metrics	
Monitoring and administering applications, departments, and general groups	
Applications	
Departments	
Application and department hierarchies	
Exporting information about applications and departments	
General groupsProvisioning storage	
Configuring IBM Spectrum Control for provisioning	
Provisioning storage with the IBM Spectrum Control GUI	
Optimizing storage tiering	
Investigating the capacity of tiered storage	440
Setting the tier level of storage pools	
Renaming tiers	
Analyzing tiering by servers	
Analyzing tiering by hypervisors	
Analyzing tiering by storage virtualizers	
Analyzing tiering by storage pools	
Analyzing tiering by volumes	
Tiering volumes by I/O density and I/O rate	
Modifying the criteria for analyzing tiering	
Optimizing storage pools	
Balancing pools	
Modifying the criteria for balancing pools	
Criteria for identifying the pools that require balancing	
Reclaiming storage	
Viewing storage reclamation	
Viewing storage reclamation by storage system	

Excluding volumes from reclamation analysis	
Transforming and migrating volumes	
Modifying the period for analyzing performance data	449
Modifying the properties of resources	
Properties of resources that can be modified	450
Opening the management GUI for a resource	452
Opening the DS Storage Manager GUI for a DS8000	452
Exporting information to a file	452
Customizing lists	454
Filtering lists	454
Sorting lists	455
Showing, hiding, and reordering columns in lists	455
Managing tasks	455
Viewing tasks	456
Managing tasks for provisioning	459
Managing tasks for tiering storage, balancing pools, and transforming storage	461
Renaming tasks	466
Deleting tasks	466
Planning copy data resources	467
Monitoring copy data resources	467
REST APIs	468
Retrieve data by using a REST API command line utility	468
Retrieve data by using REST APIs with a web browser	470
Chapter 7. Reporting	473
Creating predefined capacity reports	475
Tutorial: Creating a predefined capacity report about storage systems	478
Tutorial: Creating a predefined capacity report about pools	479
Tutorial: Creating a predefined capacity report about the volumes assigned to servers	480
Tutorial: Creating a predefined capacity report about managed disks by storage systems.	481
Tutorial: Creating a predefined capacity report about tiered pools	481
Creating predefined inventory reports	
Tutorial: Creating an inventory report about block storage systems	
Creating custom reports	
Tutorials: Creating custom capacity and performance reports for applications	
Custom capacity and performance view reports	
Creating chargeback and consumer reports	
Creating chargeback reports	
Creating consumer reports	
Creating summary reports of the storage capacity	
Running reports	
Editing reports	
Deleting reports	
Configuring the email server	
Investigating issues with reports	
Types of predefined capacity and inventory reports	
Reports FAQ	503
Exporting information about resources	504
Exporting information about capacity shortages in pools	
Exporting information about the input/output performance of volumes	
Capacity metrics for reports	
Capacity metrics for chargeback and consumer reports	506
Appendix A. Reference	509
Opening IBM Spectrum Control on Windows operating systems	
Opening IBM Spectrum Control GUIs and CLIs	
Accessing administration tools	510

Appendix B. Accessibility features for IBM Spectrum Control	513
Legal notices	517
Privacy policy considerations	518
Trademarks	519
Glossary	521
A	
В	
C	
D	
E	
F	
G	
Н	
I	529
J	
K	530
L	
M	530
N	531
0	531
P	532
R	533
S	534
Т	537
U	537
V	538
W	539
Z	539
Indax	E / /

About this guide

IBM Spectrum Control manages storage infrastructure by centralizing, automating, and simplifying the management of complex and heterogeneous storage environments. This guide provides task-oriented user information that helps you to obtain optimal product performance.

Who should read this guide

This publication is intended for administrators and users who want to learn how to use IBM Spectrum Control to monitor and manage storage resources and infrastructure.

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

- Storage Area Network (SAN)
- Simple Network Management Protocol (SNMP)
- Storage provisioning, optimization, performance, tiers, and cloud configuration
- IBM® DB2®
- Jazz[®] for Service Management
- IBM Tivoli® Common Reporting

Publications

A number of publications are provided with IBM Spectrum Control.

The following section describes how to access these publications online.

Accessing publications online

Information about installing, configuring, using, upgrading, and uninstalling IBM Spectrum Control and related products is available online.

Use <u>Table 1 on page ix</u> to view and download these publications. Translated documents are available for some products.

Table 1. Locations of publications for IBM Spectrum Control and related products		
Product	Online location	
IBM Cognos® Analytics	https://www.ibm.com/support/knowledgecenter/SSEP7J	
IBM® Db2® Database	https://www.ibm.com/support/knowledgecenter/SSEPGG	
IBM FlashSystem® 5000, IBM FlashSystem 5100, and IBM Storwize® V5000E	https://www.ibm.com/support/knowledgecenter/STHGUJ	
IBM FlashSystem 7200 and IBM Storwize V7000	https://www.ibm.com/support/knowledgecenter/ST3FR7	
IBM FlashSystem V840 Enterprise Performance Solution	http://www.ibm.com/support/knowledgecenter/ST2HTZ/	
IBM FlashSystem 840 and IBM FlashSystem 900	http://www.ibm.com/support/knowledgecenter/STKMQB_1.6.1	
IBM FlashSystem 9200 and IBM FlashSystem 9100	http://www.ibm.com/support/knowledgecenter/STSLR9/	

Table 1. Locations of publications for IBM Spectrum Control and related products (continued)		
Product Online location		
IBM FlashSystem V9000	http://www.ibm.com/support/knowledgecenter/STKMQV	
IBM FlashSystem A9000	http://www.ibm.com/support/knowledgecenter/STJKMM	
IBM FlashSystem A9000R	http://www.ibm.com/support/knowledgecenter/STJKN5	
IBM Spectrum Accelerate	http://www.ibm.com/support/knowledgecenter/STZSWD/welcome	
IBM Spectrum Control	http://www.ibm.com/support/knowledgecenter/SS5R93	
IBM Storage Insights	http://www.ibm.com/support/knowledgecenter/SSQRB8	
IBM Spectrum Virtualize as Software Only	https://www.ibm.com/support/knowledgecenter/STVLF4/	
IBM Spectrum Virtualize for Public Cloud	https://www.ibm.com/support/knowledgecenter/STHLEK/	
IBM Spectrum Scale (General Parallel File System)	https://www.ibm.com/support/knowledgecenter/SSFKCN/gpfs_welcome.html	
IBM Storwize V3500	https://www.ibm.com/support/knowledgecenter/STLM6B	
IBM Storwize V3700	https://www.ibm.com/support/knowledgecenter/STLM5A	
IBM Storwize V5000, IBM Storwize V5100, and IBM Storwize V5000E	https://www.ibm.com/support/knowledgecenter/STHGUJ	
IBM Storwize V7000 Unified	https://www.ibm.com/support/knowledgecenter/ST5Q4U	
IBM System Storage® DS8000®	http://www.ibm.com/support/knowledgecenter/	
	Go to the preceding website and search for the DS8000 series system that you want.	
IBM System Storage N series	http://www.ibm.com/support/knowledgecenter/nseries/kc_welcome_nseries.html	
IBM SAN Volume Controller	http://www.ibm.com/support/knowledgecenter/STPVGU	
IBM XIV® Storage System	http://www.ibm.com/support/knowledgecenter/STJTAG	
IBM WebSphere® Application Server	https://www.ibm.com/support/knowledgecenter/SSEQTP/mapfiles/product_welcome_was.html	
IBM WebSphere Application Server Liberty	https://www.ibm.com/support/knowledgecenter/ SSEQTP_liberty/as_ditamaps/was900_welcome_liberty.html	

IBM Redbooks

The IBM Redbooks® are publications about 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.

Languages

Non-English publications are available from IBM Knowledge Center, which is available in certain non-English languages. It is displayed in the language that is appropriate for the browser locale setting.

When a locale does not have a non-English version, the information is displayed in English, which is the default language. Non-English PDFs are available when the information is translated.

Providing feedback about publications

Your feedback is important to help IBM provide the highest quality information.

To provide comments or suggestions about the product documentation, go to the IBM Knowledge Center for IBM Spectrum Control: http://www.ibm.com/support/knowledgecenter/SS5R93. Click **Feedback** at the bottom of any page and complete the form.

IBM Spectrum Control technical community

Connect, learn, and share with storage professionals: product support technical experts who provide their perspectives and expertise.

Access the IBM Spectrum Control technical community at https://www.ibm.com/developerworks/ servicemanagement/.

Use IBM Spectrum Control technical community in the following ways:

- Become involved with transparent development, an ongoing, open engagement between other users and developers of IBM products.
- Connect one-on-one with the experts to collaborate and network about IBM and the Storage Management community.
- Read blogs to benefit from the expertise and experience of others.
- Use forums to collaborate with the broader user community.

Contacting IBM Software Support

You can contact IBM Software Support by phone, and you can register for support notifications at the technical support website.

• Go to the IBM Spectrum Control technical support website at http://www.ibm.com/support/entry/ portal/product/tivoli/ibm_spectrum_control/ibm_spectrum_control_standard_edition.

To receive future support notifications, sign in under **Subscribe to support notifications**. You are required to enter your IBM ID and password. After you are authenticated, you can configure your subscription for IBM Spectrum Control technical support website updates.

- Customers in the United States can call 1-800-IBM-SERV (1-800-426-7378).
- For international customers, go to https://www.ibm.com/planetwide/.

You can also review the *IBM Software Support Handbook*, which is available 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 IBM Software Support, 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" on page xi.

Reporting a problem

Provide IBM Software Support with information about the problems that you report.

Have the following information ready when you report a problem:

- The IBM Spectrum Control 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.

Conventions used in this guide

Information is given about the conventions that are used in this publication.

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

The following typeface conventions are used in this publication:

Bold

- · Flags that display with text
- Graphical user interface (GUI) elements (except for titles of windows and dialogs)
- · Names of keys

Italic

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

monospace

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

For syntax notations, remember the following details.

- In AIX®, the prompt for the root user is #.
- In AIX and Linux®, the commands are case-sensitive, so you must type commands exactly as they are shown.

What's new for IBM Spectrum Control Version 5.4.0


New features and enhancements are available in IBM Spectrum Control 5.4.0.

What's New	Features
Enhanced support for switches and fabrics	"Connecting directly to Brocade switches and fabrics" on page xiii
More support for non-IBM devices	"Monitoring performance for Hitachi Virtual Storage Platform F and G Series" on page xiv
Enhanced support for DS8000	"Viewing the name and VOLSER properties for DS8000 CKD volumes" on page xiv
Enhanced support for IBM Spectrum Virtualize	"User roles for monitoring IBM Spectrum Virtualize" on page xiv
Licensing and upgrade changes	 "Name updates for licenses" on page xiv "Enhancements to the Licensing page" on page xv "Upgrade update" on page xvi

Connecting directly to Brocade switches and fabrics

You can now connect directly to Brocade switches that run Fabric OS 8.2.1 or later.

That's good news if you're planning to upgrade your Brocade switches soon, or to invest in newer switches. You no longer need Brocade Network Advisor (BNA) or extra servers to monitor your switches and fabrics. Connect directly to the switch by using the Fabric OS REST API and get the same key inventory and performance information that you got when you used BNA. Using the REST API is simple and reliable way to monitor your switches.



Monitoring performance for Hitachi Virtual Storage Platform F and G Series

You can now monitor the performance of Hitachi Virtual Storage Platform F and G Series storage systems. You can also define performance alerts so that you're notified of bandwidth, latency, and other issues before they impact your storage environment. And, you can compare your storage workloads and response times across your entire storage environment.

For example, view the Total I/O Rate and the Overall Response Time of your Hitachi VSP system.



To confirm details about the supported storage systems including minimum firmware levels, review the Hitachi support page.

See "Performance metrics for Hitachi VSP storage systems" on page 332.

Viewing the name and VOLSER properties for DS8000 CKD volumes

For DS8000, you can now view both the volume name and the volume serial number (VOLSER) for your count-key-data (CKD) volumes. On the **Volumes** page, the following changes were made:

- **Name** shows the name that was assigned to the volume when it was created. Previously, the VOLSER was shown as the name of the volume. In IBM Spectrum Control, the name is the primary identifier for the volume.
- VOLSER is a new column that shows the volume serial number for DS8000 CKD volumes.

Don't forget to update tools that use the REST API: The REST API for IBM Spectrum Control was also updated so that you can retrieve the name and VOLSER properties for the volume. Ensure that you update any 3rd-party tools or reports that you use to show information that is retrieved by the REST API.

User roles for monitoring IBM Spectrum Virtualize

You can now use a role with a lower level of authority on storage systems that run IBM Spectrum Virtualize 8.3.1.2 or later to collect performance metadata. When you add these storage systems, users with a lower level of authority than Administrator, such as users with the Monitor role, can collect performance metadata.

See "User roles for collecting performance metadata from IBM Spectrum Virtualize" on page 126.

Name updates for licenses

The following licenses were renamed:

Names of IBM Spectrum Control 5.3 licenses	Names of IBM Spectrum Control 5.4.0 licenses	
IBM Spectrum Control Standard Edition	IBM Spectrum Control	
IBM Spectrum Control Standard Select Edition	IBM Spectrum Control Select Edition	

These licensing changes are automatically applied. No action is required on your part. For example, if you owned the *IBM Spectrum Control Standard Edition* in previous releases, your license is now called *IBM Spectrum Control*.

Standard Edition bonus: If you owned a *Standard* license in a previous release, you now have access to the following features:

- Chargeback, consumer, and rollup reports
- · Storage reclamation

Advanced Edition: If you own the Advanced Edition, you are not impacted by the renaming of product licenses.

Discontinued features: For all product licenses, the support for some product features were discontinued in this release. For a complete list, see "Discontinued features in IBM Spectrum Control" on page xxx.


Enhancements to the Licensing page

The **Licensing** page (**Settings** > **Licensing**) was updated to make it easier for you to understand your licensing and determine how much storage you can add before exceeding your license.

The following updates were made:

- The automatic assignment of storage category for most storage systems. Fewer manual assignments are now required.
- Improved notifications to make you aware of actions that you must complete.
- The names of column headings were changed and the hover help was revised to help you complete your compliance check.
- The capacity to be licensed is now shown in tebibytes (TiB) for consistency with the other capacity values on the page.

Tip: For storage virtualizers and IBM Spectrum Scale, the capacity to be licensed includes the capacity from back-end storage systems that are not monitored by IBM Spectrum Control. The license for the monitored back-end storage is recorded separately and has its own record on the **Licensing** page.



See "Licensing" on page 30.

Upgrade update

You cannot upgrade from any release before IBM Spectrum Control 5.3.0. If you are currently running a version of IBM Spectrum Control earlier than 5.3.0, you must upgrade to IBM Spectrum Control 5.3.0 or later.

For more information on platform support, see FIBM Spectrum Control interoperability matrix.

Product fixes and security updates

In every release of IBM Spectrum® Control, IBM fixes issues that are reported by our internal teams and by our customers. Along with fixes, improving the operation and security of the product is also a priority. See what's been updated and fixed in this release:

- Fixes that are included in this release
- Security bulletins for IBM Spectrum Control

Limitations and known issues

For information about limitations and known issues that might occur when you use IBM Spectrum Control 5.4.0, see Limitations and known issues.

Related information

- "Change History for IBM Spectrum Control" on page xvi
- Chapter 3, "Product licenses," on page 25
- "Sponsor user program" on page xxxii
- "Beta program" on page xxxii
- "IBM Storage Insights for IBM Spectrum Control" on page 54

Change History for IBM Spectrum Control

View a history of the changes in previous versions of IBM Spectrum Control.

- "Changes in 5.3.7" on page xvii
- "Changes in 5.3.6" on page xvii
- "Changes in 5.3.5" on page xviii
- "Changes in 5.3.4" on page xix
- "Changes in 5.3.3" on page xix
- "Changes in 5.3.2" on page xx
- "Changes in 5.3.1" on page xxi
- "Changes in 5.3.0" on page xxi
- "Changes in 5.2.17" on page xxii
- "Changes in 5.2.16" on page xxii
- "Changes in 5.2.15" on page xxii
- "Changes in 5.2.14" on page xxiii
- "Changes in 5.2.13" on page xxiv
- "Changes in 5.2.12" on page xxiv
- "Changes in 5.2.11" on page xxv
- "Changes in 5.2.10" on page xxv
- "Changes in 5.2.9" on page xxvi

Changes in 5.3.7

The following features and enhancements were available in IBM Spectrum Control 5.3.7:

Setting capacity limits

Set capacity limits for storage systems and pools. You want to know how much capacity you have left before your storage systems or pools are 80% full? Just set a capacity limit and you'll know how close you are to reaching your capacity limit. Then, check the Capacity-to-Limit (GiB) value and see how much capacity you can use before you reach your limit.

Support for Pure Storage devices

Monitor Pure FlashArray//M and FlashArray//X. View information about the capacity, space usage, and performance of this non-IBM storage. Other features, such as alerting, health monitoring, advanced analytics, and reporting are also supported.

More support for IBM FlashSystem devices

You can now view information about the capacity, space usage, and performance of the FlashSystem 5000, FlashSystem 5100, FlashSystem 7200, and FlashSystem 9200 storage systems.

Monitoring DS8000 capacity on flash drives

Monitor the volume capacity that Easy Tier® places on Tier 1 and Tier 2 flash, high-capacity drives on DS8000 storage systems. View the capacity and available capacity of Tier 1 and Tier 2 flash drives in a pool and the distribution of volume extents across each of the Easy Tier drive classes.

Monitoring compliance with your license

For storage systems that use the capacity license model, view the capacity breakdown of the storage environment by the categories of drives that are used. In this release, Category 1 was added for Storage Class Memory (SCM) drives and managed disks on Spectrum Virtualize for Public Cloud.

View more information about the changes in 5.3.7

Changes in 5.3.6

The following features and enhancements were available in IBM Spectrum Control 5.3.6:

Monitoring performance for Dell EMC Unity and NetApp storage systems running ONTAP 9

Support for monitoring and alerting on performance is now added for the following non-IBM storage systems:

- Dell EMC Unity
- NetApp storage systems that are running ONTAP 9

Monitoring Hitachi Virtual Storage Platform F Series and G Series

Add Hitachi Virtual Storage Platform (VSP) F Series and G Series to the list of non-IBM storage systems that you can monitor and report on directly in IBM Spectrum Control.

More support for IBM FlashSystem devices

You can now view information about the capacity, space usage, and performance of the FlashSystem 5000, FlashSystem 5200, and FlashSystem 9200 storage systems.

Gaining insights into capacity usage and capacity growth

When you log in to IBM Spectrum Control, you want to know how much capacity you have, how much capacity you've used, and how much usable capacity is still available. To make this information more consumable, the overview charts have got a makeover. But, that's not all. Now, with a click, you can access key capacity metrics for each pool and volume in your storage systems. And, to help you plan capacity, two new metrics were added:

- Recent Fill Rate (%)
- Recent Growth (GiB)

Aligning capacity terms

To enhance your experience of our products and to provide a unified and simplified view of capacity concepts, IBM Storage is aligning the capacity terms across all IBM Storage products. To ensure that

we use the same capacity term for the same capacity concept across our products, IBM Spectrum Control has changed 47 of the capacity terms that it uses, removed 4 terms that are no longer needed, and added 1 new term.

Viewing capacity savings for devices that support hardware compression

You can now view the amount and percentage of capacity that is saved for devices that use inline data compression technology, such as IBM FlashSystem 9100 and IBM Storwize V7000 Gen3.

Monitoring the cache fullness of pools on IBM Spectrum Virtualize

See how full your cache is by adding new cache fullness metrics to the performance charts when you view the performance of pools in your storage systems. Add and track cache fullness metrics to identify the pools that are experiencing heavy cache usage. You can also use the metrics to help investigate problems with volumes in a pool and to determine whether to move a volume to a different I/O group where the pool's cache partition does not have a cache fullness problem.

Monitoring the aggregated cache fullness of nodes at storage system level on IBM Spectrum Virtualize

See how full your nodes cache is at storage system level by adding new cache fullness metrics to the performance charts when you view the performance of storage systems. Add and track cache fullness metrics to identify the nodes that are experiencing heavy cache usage. You can also use the metrics to help investigate problems for volumes in a pool and to determine whether to move a volume to a different I/O group where the pool's cache partition does not have a cache fullness problem.

Monitoring unmap operations at host connection level on IBM Spectrum Virtualize

Track the workload of unmap operations that each host contributes to the system. To do this, you can measure the performance of unmap volume operations at host connection level. Metrics, such as the average number of MiBs that are unmapped from volumes, I/O rates, data rates, and response times for unmap operations, are collected and shown.

Monitoring Storage Class Memory (SCM) on IBM Spectrum Virtualize

You can now monitor the volume capacity that Easy Tier® places on SCM drives on IBM Spectrum Virtualize systems, such as IBM FlashSystem 9100, IBM FlashSystem 7200, and the IBM Storwize family.

Creating inventory reports about IP ports

You can now create and schedule predefined reports about the IP ports on the nodes in IBM Spectrum Virtualize systems, such as SAN Volume Controller, IBM FlashSystem 9100, FlashSystem V9000, FlashSystem V840, and the IBM Storwize family.

Adding IP ports to general groups

You can now add IP ports to general groups so that you can receive alert notifications about changes in the configuration or attributes of the IP ports in a general group.

Monitoring disk utilization for RAID arrays on DS8000 storage systems

For RAID arrays on DS8000 version 8.5.0 and later, the disk utilization values were updated in performance charts and reports. You can now more accurately check how busy the disks in the array are over a period and identify arrays that are underutilized and overutilized.

Support for IBM Cognos Analytics

IBM Spectrum Control now supports IBM Cognos Analytics Version 11.1.4. Note that IBM Cognos Analytics *does not* run on Windows Server 2019 and Red Hat Enterprise Linux 8.

View more information about the changes in 5.3.6

Changes in 5.3.5

The following features and enhancements were available in IBM Spectrum Control 5.3.5:

Important capacity updates

Capacity values were updated to better reflect capacity usage in devices that support hardware compression. Capacity, such as pool capacity and allocated space, is now measured based on the physical capacity of the device, rather than the effective capacity.

The devices affected by this update are:

- Storage systems such as FlashSystem 9100, FlashSystem 900, and Storwize V7000 Gen3, which contain IBM FlashCore® Modules with hardware compression.
- Storage virtualizers such as SAN Volume Controller and Storwize family that use back-end storage systems with hardware compression.

The effective capacity information is still available. Additional columns were added to the **Block Storage Systems** and **Pools** pages to show the effective capacity based on the data reduction savings that are achieved with hardware compression. Check the column names that begin with "Effective", for example, Effective Capacity (GiB) and Effective Used Capacity (%).

This update also affects alert thresholds and historical charting for storage systems and pools. In historical capacity charts, you'll see a sharp drop in the values for capacity and used space.

Learn more about the capacity updates.

Product fixes and security updates

In every release of IBM Spectrum Control, IBM fixes issues that are reported by our internal teams and by our customers. Along with fixes, improving the operation and security of the product is also a priority. See what's been updated and fixed in this release:

- Fixes that are included in this release
- Security bulletins for IBM Spectrum Control

View more information about the changes in 5.3.5

Changes in 5.3.4

The following features and enhancements were available in IBM Spectrum Control 5.3.4:

Monitoring unmap operations on IBM Spectrum Virtualize

You can now measure the performance of unmap volume commands and define alerts so that you are notified when unmap operations values fall outside thresholds you specify. Metrics, such as the average number of MiBs that are unmapped from volumes, I/O rates, data rates, and response times for unmap operations, are collected and shown:

- At the pool level
- At the node level
- At the I/O group level

For more information, see "Performance metrics for resources that run IBM Spectrum Virtualize" on page 289.

Enhanced OS support

You can now install IBM Spectrum Control and Storage Resource agents on servers that run Windows Server 2019 and Red Hat Enterprise Linux 8.

Product fixes and security updates

In every release of IBM Spectrum Control, IBM fixes issues that are reported by our internal teams and by our customers. Along with fixes, improving the operation and security of the product is also a priority. See what's been updated and fixed in this release:

- Fixes that are included in this release
- Security bulletins for IBM Spectrum Control

View more information about the changes in 5.3.4

Changes in 5.3.3

The following features and enhancements were available in IBM Spectrum Control 5.3.3:

IBM Storage Insights for IBM Spectrum Control

IBM Storage Insights for IBM Spectrum Control is an IBM Cloud service that can help you predict and prevent storage problems before they impact your business. It is complementary to IBM Spectrum

Control and is available at no additional cost if you have an active license with a current subscription and support agreement for IBM Virtual Storage Center, IBM Spectrum Storage Suite, or any edition of IBM Spectrum Control.

For more information, see IBM Storage Insights for IBM Spectrum Control.

Monitoring IBM Storwize V5000E storage systems

You can now monitor Storwize V5000E storage systems. You can view information about the capacity, space usage, and performance of the storage systems. Other features, such as alerting, health monitoring, advanced analytics, and reporting are also supported.

Viewing application performance

You can now view the total workload of an application that is consuming storage resources. Previously, to investigate the workload of an application, you had to review the workloads of the individual volumes. Now you can select an application from the **Applications page** and click **View Performance** to see the **Total I/O Rate** and the **Total Data Rate** for the application. On the applications details page, you can view the **Aggregated Workload** charts.

Faster ways to view and acknowledge your alerts

Now it's easier and faster to view and acknowledge your alerts. On the **Alerts** dashboard, just double-click an alert to see details of the alert on the same screen, including a performance chart for performance alerts. You can view summary details of several alerts at the same time, and acknowledge them all with a single click.

To filter alerts, you can now click one of the alert category tags. For example, to show critical and warning unacknowledged alerts, click **Critical** and **Warning**. To show your acknowledged alerts, click **Acknowledged**.

Automate changing the password for Spectrum Control on Windows

You can now run the changepassword. bat script from Windows Command Line Interface, or from an automated script, to change the passwords for the Spectrum Control user ID or for the database user ID.

Improved support for special characters in Spectrum Control and DB2 passwords

IBM Spectrum Control has improved the support for special characters in passwords, giving you the flexibility that you need based on your security requirements.

IBM Spectrum Control ports have TLS 1.1 and 1.0 disabled by default to increase security

Prior to version 5.3.3, IBM Spectrum Control ports had TLS 1.1 and 1.0 enabled by default. To increase security IBM Spectrum Control ports have TLS 1.1 and 1.0 disabled by default. Therefore, IBM Spectrum Control won't be able to communicate with devices that do not support TLS 1.2. If you want to upgrade your devices to a version that supports TLS 1.2, contact your vendor.

View more information about the changes in 5.3.3

Changes in 5.3.2

The following features and enhancements were available in IBM Spectrum Control 5.3.2:

Configuring alerts using alert policies

Configuring alerts is now easier! You can now use alert policies to configure alerts for groups of resources rather than just for individual resources.

Simply create an alert policy and specify the alerts that you want, and those alerts are automatically applied to all the resources that you include in the policy. Modifying alerts for resources is also a snap because the changes that you make are automatically applied to all the resources in the policy at the same time.

For more information about alerts and alert policies, see "How alerts work" on page 152.

Reporting

- Create and schedule predefined reports about the capacity that is assigned to hypervisors.
- Create an inventory report about your servers' ports showing information about the ports that your servers are connected to.

- Easily generate predefined capacity reports about the resources that you are interested in. It just takes a click to select managed disks by storage systems, pools by storage systems, servers by applications, volumes by servers.
- Send custom and predefined reports as CSV and HTML attachments. Recipients can then download and open the report in a tool of their choice.
- Can't wait for the next run of a monthly report? Just click and run the report now to share the information with your colleagues!

For more information about reporting, see Chapter 7, "Reporting," on page 473.

Monitoring compression savings for RAID arrays in IBM FlashSystem 900, model AE3 storage system For FlashSystem 900, model AE3 with firmware version 1.6, you can now view the amount of capacity that is saved when the flash modules in the RAID array use inline data compression.

Identifying mirrored volume relationships on storage systems that run IBM Spectrum Virtualize

To help you identify the primary and secondary copies in mirrored volume relationships, the following properties for volumes were renamed:

- Target is now named Mirror Role and shows the value "Primary" or "Secondary".
- Mirror is now named Copy ID.

Tip: The name changes for these columns are also reflected in the volume information that is retrieved by using the REST API for IBM Spectrum Control. Ensure that you update any 3rd-party tools or reports that you use to show information from the REST API.

View more information about the changes in 5.3.2

Changes in 5.3.1

The following features and enhancements were available in IBM Spectrum Control 5.3.1:

- Reporting enhancements:
 - From a single page, create predefined capacity reports that you can schedule, send by email and save to your file system.
 - Create predefined inventory reports for your storage systems and for the storage resources that they
 use, such as nodes.
 - Add filters to custom reports. For example, you want to create a report that shows whether your file systems have enough capacity.
- You can now monitor Storwize V7000 Gen3 storage systems. You can view information about the capacity, space usage, and performance of the storage systems. You can also view the savings that are achieved when data is reduced, compressed, and deduplicated.

View more information about the changes in 5.3.1

Changes in 5.3.0

The following features and enhancements were available in IBM Spectrum Control 5.3.0:

- Create custom reports. From any table view, create reports to share information with your colleagues and managers about the capacity, configuration, and health status and the performance of the resources that you monitor in your storage environment.
- Monitor IBM FlashSystem 9100 storage systems.
- Prevent your FlashSystem 900 from running out of capacity by monitoring the physical capacity and receiving alerts when the physical capacity is low. View information about pool capacity and the used and available capacity in the RAID array.
- Add multiple storage systems for monitoring at the same time.

View more information about the changes in 5.3.0

Changes in 5.2.17

The following features and enhancements were available in IBM Spectrum Control 5.2.17:

- View detailed information about the enclosures for block storage systems that run IBM Spectrum Virtualize, such as FlashSystem V9000, FlashSystem V840, FlashSystem 900, SAN Volume Controller, and the IBM Storwize family.
- Use new overview charts for storage systems to answer key physical capacity, volume capacity, and capacity savings questions:
 - How much space is available.
 - How much of your thin-provisioned volume space is used.
 - How much space you are saving by using compression, deduplication, and thin provisioning.
- For storage systems that run IBM Spectrum Virtualize version 8.1.1 or later, use the Pool Attributes column to distinguish data reduction pools from standard pools and identify if the data in the pool is encrypted. For storage systems that run IBM Spectrum Virtualize version 8.1.3 or later, you can also view the capacity savings that are made when you use data reduction pools.
- For storage systems that run IBM Spectrum Virtualize version 8.1.3 or later, you can see the amount and percentage of capacity that is saved when data is deduplicated. The capacity savings for deduplication are shown at the storage systems, pools, and volumes levels.
- For storage systems that run IBM Spectrum Virtualize, you can track the historical capacity and the capacity savings that were achieved when volumes are compressed. This allows you to monitor the efficiency of compression over time.
- Use the Compressed and Deduplicated filters to show only volumes with the selected filter value.

View more information about the changes in 5.2.17

Changes in 5.2.16

The following features and enhancements were available in IBM Spectrum Control 5.2.16:

- Monitor the availability of clusters by checking which nodes have spare nodes. You can also check the status of the spare nodes, such as which nodes are online and which nodes are on standby.
- Monitor the workload of compressed volumes on storage systems that run IBM Spectrum Virtualize[™].
 You can view I/O rate, data rate, and response time metrics for compressed volumes and uncompressed volumes. See Compression metrics for volumes.
- In previous releases, IBM Cognos Business Intelligence reports was installed with and could be accessed directly from IBM Spectrum Control. In this release, IBM Cognos Analytics is installed as a separate product. See Migrating reports to IBM Cognos Analytics during upgrade.
- The integration with external LDAP repositories in IBM Spectrum Control is improved.

View more information about the changes in 5.2.16

Changes in 5.2.15

The following features and enhancements were available in IBM Spectrum Control 5.2.15:

- Creating agentless servers used to be a manual process, and automating it was a frequently requested enhancement. IBM Spectrum Control now automatically generates agentless servers for the physical servers and virtual machines in your SAN environment, so that you can easily see how they are consuming storage.
- Change the default names of the tiers that are shown in the GUI.
- Add or change properties, such as the location or custom tags, for multiple servers, storage systems, hypervisors, switches, or fabrics. You can use the properties to filter or sort the resources in the GUI or in an external application if the data is shared or exported.
- Define more than 30 new alerts for volumes in block storage systems, including alerts for capacity values such as Compression Savings, Data Deduplication Savings, Data Reduction Savings, and

Unallocated Volume Space. Use these alerts to detect and monitor potential configuration and capacity issues in your environment.

- Find volumes with the same attributes. Use the new filters, such as Easy Tier and Mirror, to show only volumes with the selected filtered value. You can also sort the volumes by allocated and unallocated space, and much more.
- Add and monitor IBM Spectrum Virtualize software-only clusters.
- IBM DB2 Version 11.1 is now included with IBM Spectrum Control.
- Use the Representational State Transfer (REST) API to automatically export the data that you need to generate custom reports or share with external applications. The following information is now available:
 - How much capacity is used by applications, departments, hypervisors, and physical servers.
 - How much block capacity is used by a specific application, department, hypervisor, or physical server.

View more information about the changes in 5.2.15

Changes in 5.2.14

The following features and enhancements were available in IBM Spectrum Control 5.2.14:

- Enhancements for IBM Cloud Object Storage:
 - Add COS Vaults to applications and general groups so you can track how much space is being consumed by logically related vaults and how that consumption trends over time.
 - Use the enhanced Vaults by Accessers chart, and the new Vaults by Drives and Slicestors by Drives charts, to monitor the availability of all your COS Vaults and COS Slicestor® nodes.
 - View the number of failed drives in the COS Slicestor nodes and the tolerance of the nodes to more drive failures.
 - View the tolerance of COS Vaults to drive failures across all the COS Slicestor nodes in the storage pool.
 - View the COS Vaults, access pools, and storage pools that are used by IBM Spectrum Scale storage systems.
- The user experience is enhanced within the GUI for integration with an external LDAP repository. IBM Spectrum Control is now available to a larger set of users and groups and you are able to log in to IBM Spectrum Control with one set of credentials.
- IBM Spectrum Control ensures, when it sends alert notifications and reports by email to your email server, that it complies with the security standards that you configure on your email server.
- You can now choose the unit of measurement for a capacity metric when you create or edit chargeback and consumer reports. So, you can generate reports that show storage consumption in the unit of measurement of your choice whether it is PB or PiB, TB or TiB, or GB or GiB.
- View the following port delay metrics for storage systems that run IBM Spectrum Virtualize: Port Delay Time, Port Delay I/O Percentage
- DS8000 storage systems, you can now see the hosts and host clusters that host connections belong to.
- View a list of all the primary volumes that are not protected by a copy data relationship.
- View which Managed Disks are active Quorum disks.
- View the site name of each node in a stretched cluster.
- For Storwize storage systems, version 7.8 or later, you can now view information about flash drives by tier in storage pools and volumes. For storage pools, you can view the capacity and available capacity for solid-state drives (Tier 0) and for read-intensive flash storage (Tier 1). For volumes, you can view the capacity for solid-state drives (Tier 0) and for read-intensive flash storage (Tier 1).
- Upload logs automatically to IBM Software Support to help resolve problems quickly.

View more information about the changes in 5.2.14

Changes in 5.2.13

The following features and enhancements were available in IBM Spectrum Control 5.2.13:

- View more information about the capacity, space usage, and performance of Dell EMC. Support for versions of Dell EMC SMI-S Provider or Dell EMC Solutions Enabler that are compliant with SMI-S 1.6 provides a more representative view of Dell EMC VMAX, Dell EMC VNX, and Dell EMC VNXe storage systems.
- Gain insights into the external storage that is used by IBM Spectrum Scale:
 - View capacity values for the external storage that clusters and file systems use. This view provides
 insight into the total amount of data that is migrated from all file systems in a storage system, or the
 amount of data that is migrated from specific file systems.
 - Check whether you have enough internal space to recall data that was migrated to external storage.
- For storage systems that run IBM Spectrum Virtualize, you can now see the cluster that host connections belong to.
- Enhancements for IBM Cloud Object Storage:
 - View information about the hard disk drives, solid-state drives, and flash modules in COS Slicestor nodes.
 - In the Vaults by Accessers chart, view the vaults that cannot be accessed and the vaults that are at risk of access failure.
- On the new VDisk Mirrors and HyperSwap® pages in the Copy Data panel, you can analyze storage device relationships across your block storage environment in IBM Spectrum Control to identify how redundancy is affecting available capacity.
- You exclude the cost of storage in chargeback and consumer reports when you want to highlight storage consumption and not storage cost.

View more information about the changes in 5.2.13

Changes in 5.2.12

The following features and enhancements were available in IBM Spectrum Control 5.2.12:

- Monitor the capacity, space usage, and performance of FlashSystem A9000 and FlashSystem A9000R storage systems.
- Create consumer reports that show the capacity and the cost of the block storage that an application, department, hypervisor, or physical server uses.
- Report on allocated versus assigned capacity in chargeback and consumer reports.
- · Analyze replication and FlashCopy relationships across your block storage environment
- View the amount of external storage from cloud services and other storage providers that is being used by the file systems in a IBM Spectrum Scale cluster. For external storage that is provided by IBM Cloud Object Storage, you can also view capacity information, including the percentage of space that is being used.
- Monitor sites, mirrors, and vault quotas for IBM Cloud Object Storage. For example, you can:
 - Monitor the capacity and space usage of sites. See which COS Slicestor and COS Accesser® nodes are located at each site and which sites cannot tolerate COS Accesser node failures.
 - View the space quotas that are configured for vaults and see whether any vaults are violating the soft or hard quota limits.
 - View information about mirrors, such as the pair of vaults in each of the mirrors, and the number of COS Accesser nodes that are configured to access the mirrors. You can also see the storage pools and sites for the vaults in the mirror.
- Define alerts that notify you of status changes on the IBM Spectrum Scale nodes that are configured as cloud gateways.

- Customize the analysis period for analyzing tiering and balancing pools to the days of the week and hours of the day of your choice.
- Access information about your resources by using the Representational State Transfer (REST) API for IBM Spectrum Control. You can use this information to generate custom configuration and performance reports.

View more information about the changes in 5.2.12

Changes in 5.2.11

The following features and enhancements were available in IBM Spectrum Control 5.2.11:

- Changes to alerts:
 - Create alerts for a set of storage resources that are assigned to an application.
 - Define alerts for a set of resources in custom groups that you create.
 - Combine multiple alert conditions from multiple resources in a single, custom alert.
- Create, schedule, and send chargeback reports by email to make the owners and managers aware of the cost of the capacity and the amount of capacity that is used by their applications, departments, hypervisors, and servers.
- View key performance metrics for IBM® FlashSystem™ 840 and IBM FlashSystem 900 storage systems:
 - Response times
 - I/O rates and data rates
 - The condition of the storage system
- Review new capacity charts to detect capacity shortages and investigate trends in storage growth for your tiers.
- Analyze replication relationships across your block storage environment.
- Monitor the status, capacity, and space usage of IBM Cloud Object Storage.
- Identify which file systems in a Spectrum Scale cluster use external storage that is provided by cloud services or other storage providers.
- Enable automated probe scheduling to simplify the process of scheduling probes and help avoid excessive load on the product server.
- Determine whether the total capacity of the storage systems that IBM Spectrum Control manages exceeds your product license entitlement.
- When you complete a fresh installation of the product, the number of random listening ports is reduced. Information about the ports is documented and you can change the Device server ports after you complete the installation or perform an upgrade.
- Manually uninstall IBM Tivoli® Storage Productivity Center for Replication from IBM Spectrum Control servers on your computer.
- IBM DB2® Version 10.5 Fix Pack 7 is now included with IBM Spectrum Control.

View more information about the changes in 5.2.11

Changes in 5.2.10

The following features and enhancements were available in IBM Spectrum Control 5.2.10:

- Identify key performance metrics that are outside of a standard range on IBM System Storage® SAN Volume Controller and IBM Storwize® storage systems.
- Export performance data for a resource to a compressed file.
- Review new capacity charts to detect capacity shortages and investigate trends in storage growth for your block, file, and object storage resources.

- Review the capacity and space usage that the tiered storage in your environment consumes. You can see when your tiered storage will run out of space and the weekly growth rate in storage usage for each tier.
- View the values for the capacity of virtualized storage and non-virtualized storage to understand how capacity is used when storage virtualization is implemented in your storage environment.
- Use rollup reporting to combine capacity and status information from multiple instances of IBM
 Spectrum Control™. You can use this information to gain a network-wide perspective of storage usage in your environment when you have multiple IBM Spectrum Control servers deployed.
- Group resources, such as the storage systems with lease agreements that end in the current year, so that you can view information about the resources at one location.
- IBM Spectrum Control V5.2.10 now supports only Storage Management Initiative (SMI) data sources, also called CIM agents, for managing switches from Brocade Communications Systems, Inc.
- View the resource types that were explicitly added to an application or subcomponent.

View more information about the changes in 5.2.10

Changes in 5.2.9

The following features and enhancements were available in IBM Spectrum Control 5.2.9:

- Reclaim unused storage capacity and more effectively allocate volume space on your block storage systems.
- View the amount of storage that can be reclaimed in each of the storage systems that are monitored.
- Use the volume reclamation recommendations to reclaim unused storage space and use your storage more efficiently. By reclaiming the volumes, you can reduce costs by recycling existing storage space instead of purchasing new storage media.
- IBM Spectrum Control[™] now displays the switch blade number, port number, and blade names just as they were defined in other tools, such as Brocade Network Advisor.
- IBM DB2[®] Version 10.5 Fix Pack 6 is now included with IBM Spectrum Control.

View more information about the changes in 5.2.9

Changes in 5.2.8

The following features and enhancements were available in IBM Spectrum Control 5.2.8:

- The name of the product was changed from IBM Tivoli Storage Productivity Center to IBM Spectrum Control.
- Comprehensive improvements to alerting capabilities:
 - You can now define alerts for all the key attributes of a resource, including attributes for status, configuration, capacity, and performance.
 - To avoid receiving too many emails and disruptive notifications for some alerts, you can now choose
 a new suppression option where only one notification is sent.
 - If you want to define more than one alert for the same attribute but with different criteria and settings, you can duplicate that alert.
- View the performance of clusters, nodes, and file systems in IBM Spectrum Scale.
- Monitor object storage systems on IBM Spectrum Scale.
- Monitor the capacity and space usage of IBM FlashSystem 840 and IBM FlashSystem 900 storage systems.
- View charts that show the current and historical capacity of your block storage pools.
- IBM Spectrum Control is now a native, 64-bit application for all server operating systems.
- Simple Network Management Protocol (SNMP) Version 3 is now the preferred protocol for the management of switches and fabrics from Cisco Systems, Inc., and is an optional protocol for devices from Brocade Communications Systems, Inc. and other vendors.

- View information about the performance of inter-switch connections.
- View information about Brocade switches that are in Access Gateway mode. You can also view information about the routes from ports that are connected to a hypervisor, server, or storage system to ports that are connected to fabrics.
- View information about Brocade F_port trunks and about NPV links.
- You can now more easily identify the internal resources that are encountering problems in your environment.
- Compare how much capacity is managed by IBM Spectrum Control with the total amount of capacity that you are entitled to manage.
- View information about the capacity for each of the class drives (SSD, Nearline HDD, and Enterprise HDD), that Easy Tier can use to tier or distribute volume extents.
- To consolidate the functions of IBM Spectrum Control into a single GUI, the stand-alone GUI was removed.

View more information about the changes in 5.2.8

Capacity terminology changes

IBM Storage is making the capacity terminology that is used in IBM Storage products consistent. In Q1 2020, the capacity terminology was made clearer and simpler across IBM Storage products, including IBM Spectrum Control.

The following table lists the original capacity terms and the new terms:

Original terms	New terms	
(No existing term)	Overhead Capacity	
Allocated SAN Space	Used SAN Capacity	
Allocated Space	Used Capacity, only on Block Storage Systems, Volumes, and Pools pages	
Assigned SAN Space	Mapped SAN Capacity	
Assigned Volume Space	Mapped Capacity	
Available Disk Space	Available Drive Capacity	
Available File System Space	Available File System Capacity	
Available Pool Space	Available Capacity	
Available Repository Space	Available Repository Capacity	
Available Soft Space	Available Soft Capacity	
Available Space	Available Capacity	
Disk Capacity	Capacity	
Disk Space from name_storage_system	Drive Capacity from name_storage_system	
Effective Available Capacity	Available Written Capacity	
Effective Capacity	Written Capacity Limit	
Effective Unallocated Volume Space	Available Volume Capacity	
External Pool Used Space	External Pool Used Capacity	
External Used Space	External Used Capacity	
Effective Used Capacity	Used Written Capacity	

Original terms	New terms
Enterprise HDD Available Space	Enterprise HDD Available Capacity
File System Space from name_storage_system	File System Capacity from name_storage_system
Inactive Used Space	Inactive Used Capacity
Nearline HDD Available Space	Nearline HDD Available Capacity
Overhead Space	(Term removed)
Physical Allocation	Used Capacity, only on Block Storage Systems, Volumes, and Pools pages
Pool Capacity	Capacity
Raw Disk Capacity	Raw Capacity
Reserved Pool Space	Reserved Capacity
Soft Space	Soft Capacity
Space Quota	Capacity Quota
Tier 0 Flash Available Space	Tier 0 Flash Available Capacity
Tier 1 Flash Available Space	Tier 1 Flash Available Capacity
Total Capacity	Capacity
Total Data Reduction Savings	Total Capacity Savings
Total Disk Capacity	Capacity
Total Disk Space	Drive Capacity
Total File System Capacity	File System Capacity
Total Space	Capacity
Total Volume Capacity	Provisioned Capacity
Unallocatable Volume Space	Overprovisioned Capacity
Unallocated Space	Available Capacity
Unallocated Volume Space	Available Volume Capacity
Unassigned Volume Space	Unmapped Capacity
Unreserved Pool Space	Unreserved Capacity
Unused Space	Reserved Volume Capacity
Used Allocated Space	(Term removed, use Used Capacity instead)
Used File System Space	Used File System Capacity
Used Pool Space	(Term removed, use Used Capacity instead)
Used Space	(Term removed, use Used Capacity instead)
Virtual Allocation	Provisioned Capacity
Volume Capacity	Capacity
Written Space	Written Capacity

Restriction: The capacity term changes do not apply to the following functions:

- REST API, with the exception of the following attributes that were removed:
 - Used Space on the Block Storage Systems, Pools, and Volumes pages
 - Used Pool Space on the Storage Systems page
 - Used Allocated Space (%) on the Volumes page
- Reports in the Cognos Analytics reporting tool
- · Command-line interface

Descriptions of new capacity terms

available capacity

The amount of usable capacity that is not yet used in a system, pool, array, or managed disk.

overhead capacity

The amount of usable capacity that is occupied by metadata in a pool or system and other data that is used for system operation.

overprovisioned capacity

The extra capacity when there is more provisioned capacity in a storage system or pool than there is usable capacity. This is done when thin-provisioning or compression ensure that the used capacity of the provisioned volumes is less than their provisioned capacity.

provisioned capacity

The total capacity of all volumes in a pool or system.

raw capacity

The reported capacity of the drives in the system before formatting or RAID is applied.

reserved volume capacity

The amount of pool capacity that is reserved but has not been used yet to store data on the thin-provisioned volume.

total capacity savings

The total amount of usable capacity that is saved in a pool, system, or volume through the application of a compression or deduplication algorithm on the written data. This saved capacity is the difference between the written capacity and the used capacity.

unused capacity

The difference between the provisioned capacity of the volume and the used capacity of the volume.

usable capacity

The amount of capacity that is available for storing data on a system, pool, array, or managed disk after formatting and RAID techniques are applied.

used capacity

The amount of provisioned capacity that is taken up by data after data reduction techniques have been applied.

The amount of usable capacity that is taken up by data in a system, pool, array, or managed disk after data reduction techniques have been applied.

written capacity

The amount of provisioned capacity that would have been used to store written data on the thin-provisioned volume if data reduction techniques were not applied.

The amount of usable capacity that would have been used to store written data in a pool or system if data reduction was not applied.

written capacity limit

The amount of provisioned capacity that can be created in a system or pool without running out of usable capacity given the current total capacity savings being achieved. This capacity equals the usable capacity divided by the total capacity savings percentage. In some storage systems, restrictions in the system determine the maximum provisioned capacity allowed in a pool or system. In those cases, the effective capacity cannot exceed this limit.

Discontinued features in IBM Spectrum Control

View a list of features that were discontinued in different releases of IBM Spectrum Control.

What is a discontinued feature?: Discontinued features are features that are no longer supported in IBM Spectrum Control. In some cases, these features were removed from the GUI and are no longer available. In other cases, these features might still appear in the GUI, but they are no longer supported and might not work as intended. For both cases, it's recommended that you use another feature or product offering to complete the same actions, when possible.

Feature	Discontinued in version	Resolution
Operating systems where you can install Storage Resource agents: • HP-UX 11i v3 • Oracle Solaris 10 and 11	5.4.0	Install your Storage Resource agents on different servers with supported operating systems. For a list of supported operating systems, see https://www.ibm.com/support/docview.wss? uid=ibm10719187#Agents.
 Storage systems: DS3000, DS4000®, DS5000, or DS6000™ Non-IBM storage systems that are monitored with SMI-S 1.1 providers. NetApp devices that run ONTAP 8. 	5.4.0	No future support is planned for monitoring these storage systems. For a complete list of storage systems that are supported in 5.4.0 and later, see IBM Spectrum Control interoperability matrix for storage systems.
Provisioning of block and file storage systems, which includes: Cloud configuration concepts, such as service classes and capacity pools Reserved capacity concept External Application role	5.4.0	No other feature in IBM Spectrum Control currently replaces this feature. For your future provisioning and automation needs with IBM storage, consider using another product, such as Ansible by Red Hat. For more information about Ansible, see https://www.ansible.com/integrations/infrastructure/ibm-storage .
Optimization: • Tiering • Balancing • Volume transformation	5.4.0	No other feature in IBM Spectrum Control currently replaces this feature. It's recommended that you use a different product to optimize tiering, balance storage, and transform volumes. Pro tip: With the increasing use of all flash storage and environments with a few large pools rather than many smaller pools, optimization and balancing functions are becoming less critical.

Feature	Discontinued in version	Resolution
STAR data is no longer available in the following:	5.4.0	For your reclamation needs, go to Advanced Analytics > Reclamation.
Cognos Analytics reporting tool CLI (tpctool) commands		• For tier-related capacity management, go to Groups > Tiers .
IBM Spectrum Control GUI		For other uses of your STAR data, it's recommended that you use a different product.
Tivoli Storage Productivity Center for Replication	5.3.0	Use IBM Copy Services Manager. See IBM Copy Services Manager.
IBM Spectrum Control workflows for storage provisioning with IBM Tivoli Provisioning Manager	5.3.0	No other feature in IBM Spectrum Control currently replaces this feature.
The agent for integrating IBM Spectrum Control alerting with IBM Tivoli Monitoring	5.3.0	No other feature in IBM Spectrum Control currently replaces this feature.
OpenStack Cinder driver that uses the provisioning feature in IBM Spectrum Control	5.3.0	Use the Cinder driver that is associated with your storage system.
Database views for custom reporting with SQL	5.3.0	Use the Representational State Transfer (REST) API to access information about resources and to generate custom configuration and performance reports. See "IBM Spectrum Control REST API" on page 85.
The tpctool CLI commands for provisioning that don't use the service-class concepts of IBM Spectrum Control	5.3.0	No other feature in IBM Spectrum Control currently replaces this feature.
Using IBM Spectrum Control components to enable IBM DS8000 storage systems to work with external authentication servers, such as LDAP	5.3.0	No other feature in IBM Spectrum Control currently replaces this feature.
vSphere Web Client extension for IBM Spectrum Control and VASA 1.0 provider	5.3.0	No other feature in IBM Spectrum Control currently replaces this feature. However, IBM Spectrum Connect provides IBM Storage Enhancements for VMware vSphere Web Client and an IBM Storage Provider for VMware VASA. For more information, see the IBM Spectrum Connect Knowledge Center.

Feature	Discontinued in version	Resolution
Graphical view of the data path between storage resources, SAN components, and servers (Data Path Explorer)	5.3.0	To view the relationships between resources:
		Go to a resource list page, such as the Block Storage Systems page.
		Right-click a resource and select View Details.
		3. On the details page, view the Related Resources section. Click a related resource to view more information about it.
		Related resources that you can view include servers, file systems, fabrics, switches, object storage systems, and back-end storage systems for storage virtualizers.

Sponsor user program

Sponsor users interact directly with designers and developers to improve the user experience and to help shape the future of the overall storage portfolio.

IBM needs your experience and expertise as an active participant to work with our designers and developers to create new features for the future and help improve the overall look and feel of our IBM Storage products.

Are you ready to align us with your reality of what a product should do?

Are you a real user or a potential user of an IBM Storage software product?

IBM wants you to be excited about being an active participant and establishing a collaborative and fluid relationship between your company and our design and development teams.

Express your thoughts and get your ideas heard! Then, as the discussion moves forward, we will integrate your insights and feedback into our decision making process and make them action items. In addition, as an active program user you get a peek into upcoming storage releases and gain beta access to new products.

Note: If you are interested in joining the IBM Storage Sponsor User Program or just want to get additional information, complete the following form: IBM Storage Sponsor User.

Important: Always remember your information is kept confidential and is only used by the IBM Design and Development for product development purposes.

Beta program

The IBM Spectrum Control Beta is a continuous program. It gives you a first look at upcoming features, a chance to influence design, an opportunity to test the new features in your own environment, and a direct voice into the product development process.

There are many reasons why new and existing customers, IBM Business Partners, and even current members of IBM become part of the Beta program as a tester.

Customer quote: Bob Oesterlin, Sr. Principal Storage Engineer at Nuance: "If you're a IBM Spectrum Control user, definitely participate in IBM's Beta program. It gives you access to the latest features early

on, direct access to developers and a great community of other early adopters. I've been a part of the IBM Spectrum Control Beta program for years, and the IBM team has been fantastic and responsive."

Some benefits of the program include:

Gain access to early code and evaluate new product features and enhancements

You get access to the Beta code before general availability of the product release to determine whether the new features and enhancements are a good fit for your business organization. Once the code is downloaded, you are able to validate and test the new software directly in your environment. You can then identify and then fix any concerns before the code is available saving you precious time and preventing any production issues later on. When the code is made available, you are ready for installation, implementation and to take advantage of the new capabilities.

Interact with design and development

The product designers, architects, developers, and testers are integrated into the Beta program to help support the participants. They can assist you with any issues that you might encounter. You can also request an advocate to work with you during the process and post questions on the Beta forum.

Collaborate with other Beta customers

The Beta program includes group meetings that provide you with an opportunity to interact with other program participants about your configuration and testing experiences with the beta features. Participants are encouraged to share their experiences with the development team.

Participate in product education

The Beta program provides the participants with education on the new features and functions that are available in a Beta. The education is usually done in presentation form on a web conference where you can get a head start in learning about the capabilities.

Become an IBM reference customer

After your positive Beta experience, IBM invites you to participate in the reference program. The IBM Marketing team helps you craft a message to let other potential Beta testers know about your success when you adopt and use beta features.

Enrollment information

To enroll in the Beta program, complete the IBM Spectrum Control Beta Sign up form.

Chapter 1. Getting started

Learn about the key tasks for setting up IBM Spectrum Control for managing the resources and infrastructure in a storage environment. The getting started guide provides a list of tasks in sequential order for helping you to learn how to install, configure, and use IBM Spectrum Control.

Complete the following tasks to install and set up IBM Spectrum Control for monitoring and managing storage resources. Click the included links to learn more information about each task.

Table 2. Tasks for getting started with IBM Spectrum Control				
Task	Description	1 Learn more		
Planning for installation and upgrades.	The time that you invest in planning directly affects your ease of implementation. Planning includes the following tasks: • Ensuring that your systems meet hardware and software requirements • Becoming familiar with the operating systems, storage resources, agents, and communication protocols that can be used with IBM Spectrum Control • Ensuring that ports are open for use • Understanding authentication and rolebased authorization • Understanding the requirements for specific types of resources	Hardware requirements Software requirements. TCP/IP ports used by IBM Spectrum Control Planning for IBM Spectrum Control authentication and authorization Planning for resources For more information about the software and resources that are supported by IBM Spectrum Control, see IBM Spectrum Control interoperability matrix at http://www.ibm.com/support/docview.wss?uid=swg21386446.		

Table 2. Tasks for getting started with IBM Spectrum Control (continued)				
Task	Description	1 Learn more		
Installing or upgrading IBM Spectrum Control.	If you have a previous version of IBM Spectrum Control installed, you can upgrade it to the current version. If you do not have a previous version, you can use the following methods to install it:	You can run the installation process as a GUI wizard or through a command line in silent mode. • Description of the installation process as a GUI wizard or through a command line in silent mode.		
	Single-server environment Install IBM Spectrum Control and its components on a single server. Multiple-server environment Install IBM Spectrum Control on multiple servers if you are monitoring large storage environments, where one server is not sufficient to manage all the IBM Spectrum Control components. When you install IBM Spectrum Control, the following components are also installed: DB2 (required) IBM Spectrum Control stores the information that it collects about storage resources in a DB2 database. Jazz for Service Management and Tivoli Common Reporting (optional) Install these components if you want to use IBM Spectrum Control reports in the web-based GUI.	Installing IBM Spectrum Control in a single-server environment Installing IBM Spectrum Control in a multiple-server environment Install Jazz for Service Management and Tivoli Common Reporting		
Starting the IBM Spectrum Control GUI.	The IBM Spectrum Control GUI runs in a web browser and can be accessed from anywhere that you have a web browser and connectivity to a network.	"Starting IBM Spectrum Control" on page 44		
Adding resources for monitoring.	Before you can monitor the resources in your environment, you must add them to IBM Spectrum Control. You can add the following resources for monitoring: • Storage systems • Switches • Fabrics • Servers • Hypervisors	"Adding resources" on page 115		

Table 2. Tasks for getting started with IBM Spectrum Control (continued)		
Task	Description	Learn more
Deploying Storage Resource agents on servers (optional).	You can deploy Storage Resource agents on servers that you want to monitor. Deploy agents if you want to complete the following tasks: Collect storage usage and trending information about servers (scans) Collect database application information Discover the fabrics that are visible to a server Enforce storage policies Collect zone information	For more information about Storage Resource agents, see "Adding servers with Storage Resource agents" on page 134. For a list of the servers where you can deploy Storage Resource agents, see IBM Spectrum Control interoperability matrix at http://www.ibm.com/support/docview.wss?uid=swg21386446 and go to the Agents, Servers, and Browsers section.
Collecting data about resources.	Determining the information that you want to gather about storage resources is critical to helping you implement a storage management strategy. When you determine the types of information you want to collect about your resources, you can select the data collection jobs that you want to run. Probes Probes collect status, asset, and capacity information about resources. Performance Monitors Performance monitors collect metrics that measure the performance of storage systems and switches. Tip: Probes and performance monitors are automatically scheduled when you add resources for monitoring.	"Adding resources" on page 115 "Collecting data" on page 138
Configuring how alerts are received.	In the web-based GUI, enable alert notifications. Alerts can define notification actions that send email, generate Simple Network Management Protocol (SNMP) traps, or generate IBM Tivoli Netcool®/OMNIbus events. To receive these types of alert notifications, you must configure IBM Spectrum Control to direct the notifications to a specific email address, SNMP destination, or Tivoli Netcool/OMNIbus server.	Configuring alert notifications

Table 2. Tasks for getting started with IBM Spectrum Control (continued)		
Task	Description	① Learn more
Specifying the alerts that are generated when a problem is detected on monitored storage.	Determining when and how you are alerted to conditions or violations within your storage environment is important to helping you maintain and administer storage resources. Alerts are triggered by conditions that are detected during data collection and event processing.	 "Alert policies" on page 156 "Triggering conditions for alerts" on page 188
	To set up and manage alerts, complete the following tasks:	
	For each resource, use the default alert policy to manage the conditions and performance thresholds that can trigger alerts, and the notification settings for the alerts in the policy. Alternatively, you can create a new alert policy, or you can define alert conditions and notification settings for an individual resource.	
	View, acknowledge, and remove alerts that are detected. For alerts that are triggered by performance threshold violations, you can view a chart of the historical performance of the resource that violated a threshold.	
Configuring data retention (optional).	Specify how long to retain the data that is collected about resources and the log files that are generated by IBM Spectrum Control. By specifying the number of weeks for history retention, you can control the amount of data that is retained and available for historical analysis and charting. The longer that you retain data, the more informative your analysis, but the more storage space that is required to store that data. Default retention settings are provided.	Configuring history and data retention
Configuring authentication and user access.	Determine the authentication mechanism and authorization level for users. The authentication mechanism determines the user groups to which you can assign roles. Roles determine the IBM Spectrum Control functions that are available to the users in a group. In the web-based GUI, you can assign the following roles to groups of users: - Administrator - Monitor - External Application	Authorizing users

Managing resources

Learn about the key tasks for getting started with monitoring and managing storage resources.

After installing IBM Spectrum Control and adding resources for monitoring, complete the following tasks to manage your storage environment.

Table 3. Tasks for monitoring and managing resources		
Task	Description	1 Learn more
Monitoring the performance of resources.	Collect and view key performance metrics about resources and specify alerts to be notified of violations of performance thresholds. Use metrics to measure, identify, and troubleshoot performance issues and bottlenecks in your storage. Use alerts to identify when the performance of a monitored resource falls outside of a specified range.	"Monitoring the performance of resources" on page 260 Reporting with Cognos Analytics
Monitoring the capacity of resources.	View detailed information about the capacity of monitored resources in the GUI. You can view capacity information for storage systems, servers, hypervisors and their internal resources, such as volumes, pools, disks, and virtual machines. For example, you can view the total amount of storage capacity that is committed to a volume, the capacity that is used by a volume, and the capacity that is allocated to a volume and is not yet used. You can also use the optional Cognos Analytics reporting tool to view and create capacity reports about multiple resources. Use the information that is collected about storage capacity to help with the capacity planning of the storage within your organization.	 "Viewing information about resources" on page 231 Reporting with Cognos Analytics Click the Help icon in the web-based GUI to view descriptions of the capacity information that is displayed for resources.
Monitoring the status of resources.	Monitor the status and condition of resources to identify potential problem areas in a storage environment. On the dashboard in the IBM Spectrum Control GUI, you can view a high-level summary of the statuses for each type of resource. From the dashboard, you can drill down to view the individual status for each monitored resource.	 "How the condition of a resource is determined" on page 248 Viewing the overall status of a storage environment Viewing the status of individual resources Viewing the status of internal resources

Table 3. Tasks for monitoring and managing resources (continued)		
Task	Description	Learn more
Viewing the relationships between resources.	You can also use the optional Cognos Analytics reporting tool to view reports about the relationships between resources. For example, you can view a report about the end-to-end relationship of all the servers in your environment to back-end storage.	Viewing relationship reports in the reporting tool
Monitoring and managing applications and departments.	Storage resources can be grouped together to represent data modeling concepts such as applications, subcomponents (member applications), departments and subdepartments to create a business hierarchy. The monitoring and management of applications and departments enables you to perform capacity trending, health monitoring, and performance troubleshooting for the storage resources in your environment. In IBM Spectrum Control, you can create applications and departments, add applications and departments, use a filtering criteria to add resources to the applications, remove related resources from within the context pf applications, remove applications and departments, view key properties and view detailed storage information for applications and departments.	 "Creating applications" on page 400 "Adding subcomponents" on page 403 "Adding applications to departments" on page 409 "Using filters to assign resources to applications" on page 401 "Removing resources from applications" on page 405 "Viewing information about applications" on page 404 "Creating departments" on page 407 "Adding and creating subdepartments" on page 408 "Viewing information about departments" on page 410
Identifying hot spots of high utilization.	IBM Spectrum Control collects detailed information about storage system pools, including utilization values. The utilization value is an estimate of the average daily utilization of the physical resources such as controllers, disks, ports, and nodes that are associated with the pool. By scanning the utilization values for pools on a storage virtualizer, such as SAN Volume Controller, you can identify the hot spots or the pools that have high utilization values on the storage virtualizer. You can also determine the course of action that you take.	"Tutorial: Analyzing and re-tiering volumes in pools on tier 1" on page 88 "Tutorial: Collocating volumes" on page 90

Table 3. Tasks for monitoring and managing resources (continued)		
Task	Description	① Learn more
Setting up self- service provisioning.	Storage provisioning is the process of assigning storage to optimize the performance of a storage environment. To provision storage with IBM Spectrum Control, create service classes and capacity pools to classify and separate storage resources. After configuration, use the web-based GUI, vSphere Web Client extension, or other tool with those service classes and capacity pools to provision storage.	"Provisioning storage overview" on page 72
	In the web-based GUI, you are guided through the steps for provisioning volumes or shares. IBM Spectrum Control then recommends storage placement that is based on storage system free space and performance data. If you agree with the recommendation, IBM Spectrum Control can implement the provisioning.	
Running reports and setting up report scheduling and delivery.	Use reports to view detailed information about the condition, capacity, and performance of resources. You can use the optional Cognos Analytics reporting tool to run, schedule, and deliver predefined and custom reports. These reports include information about storage capacity, assets, and performance of storage systems, switches, fabrics, servers, hypervisors, and their internal resources.	Reporting with Cognos Analytics

Related concepts

"Managing resources" on page 111

Use IBM Spectrum Control to monitor and manage the resources and storage infrastructure in an enterprise environment. In the GUI, you can view the condition, capacity, performance, and relationships of storage resources.

Configuring and administering

Learn about the key tasks for getting started with configuring and administering IBM Spectrum Control and its related components.

Complete any of the following tasks to customize IBM Spectrum Control for your environment and administer its components.

Table 4. Tasks for configuring and administering IBM Spectrum Control		
Task	Description	C Learn more
Configuring IBM Spectrum Control.	Configure IBM Spectrum Control according to the standards within your organization, including the following tasks: • Customizing lists of resources, tasks, and alerts • Enabling reports in the optional Cognos Analytics reporting tool • Configuring how alerts are received • Configuring data retention • Configuring authentication and user access	 "Customizing lists" on page 454 If the Cognos Analytics reporting tool is not enabled, you must install Jazz for Service Management and IBM Tivoli Common Reporting. For more information, see
Administering Db2® and IBM Spectrum Control.	After IBM Spectrum Control is installed and configured, complete the following administering tasks as needed: • Administering resources and data sources • Starting and stopping IBM Spectrum Control servers • Checking the status of IBM Spectrum Control components • Packaging log files for IBM Software Support • Changing passwords for IBM Spectrum Control components • Administering Db2® and backing up the database repository	 Administering resources and data sources Starting and stopping the IBM Spectrum Control servers Checking IBM Spectrum Control status Opening PMRs and packaging IBM Spectrum Control system log files for IBM Software Support Changing passwords Administering DB2

Navigation

IBM Spectrum Control provides many of the functions for managing a storage environment. To access the functions in its GUI, use the menu bar at the top of the main window.

GUI element	Description
Page banner	Use the banner at the top of every page to complete the following tasks:
	• View the name of the product license that is active.
	 Access resource pages and main product functions from the menu bar.
	 View a list of failed tasks by using the failed tasks icon You can retry or cancel those tasks.
	View a list of tasks that are running by using the running tasks icon tasks are running.
	 View the ID and role of the user who is logged in. The role of a user determines the product functions that are available to that user. For more information about roles, see "Role-based authorization" on page 80.
	 Access the online help to view information about the currently displayed page and the overall product.
	• Enter and exit rollup mode by using the rollup mode icon icon. You must add at least one subordinate rollup server for the rollup mode icon to display in the banner.

GUI element	Description
Home	Dashboard Use the dashboard to view the overall status of monitored resources and identify potential problem areas in a storage environment. You can learn about:
	Condition and usage of resources
	Entities that consume storage on those resources
	Number and status of unacknowledged alert conditions that are detected on the monitored resources
	 Most active storage systems in your environment.
	Performance Monitors
	Use performance monitors to collect information about the performance of storage systems and switches. This information includes key performance metrics and the percentage of data collections that succeeded during the most recent 24 hours when the performance monitor was active.
	Alerts
	Use alerts to be notified when certain configuration, status, and performance conditions are detected on monitored resources.
	Tasks
	Manage the tasks that are created when you complete the following actions:
	 Place volumes on the tiers that match the workload requirements of the volume.
	Move or convert volumes.
	 Move volumes from overutilized pools to pools that are less used.
	System Management View information about the overall condition of IBM Spectrum Control. You can view information about the servers on which the product is installed. This information includes component server and database status, certain server alerts and database connection alerts, server file-system capacity information, and remote volume-performance information.

GUI element	Description
Storage	Add storage systems for monitoring and view asset, status, and performance information about their internal and related resources. You can also complete the following actions to manage and administer storage systems:
	Modify the schedules for data collection.
	Place volumes on the tiers that match the workload requirements of the volume.
	Move or convert volumes.
	Move volumes from overutilized pools to pools that are less used.
	Add resources to a capacity pool.
	Open the management GUIs for storage systems.
	Update login credentials.
	Acknowledge conditions and statuses.
Servers	Add servers and hypervisors for monitoring and view asset and status information about their internal and related resources. You can also complete the following actions to manage and administer these resources:
	Deploy Storage Resource agents to servers for full monitoring.
	Modify the schedules for data collection.
	Modify and upgrade Storage Resource agents.
	Acknowledge conditions and statuses.
Network	Add switches and fabrics for monitoring and view asset, status, and performance information about their internal and related resources. You can also complete the following actions to manage and administer these resources:
	Modify the schedules for data collection.
	Setting or modifying the zoning policy.
	Open the management GUI for a switch.
	Acknowledge conditions and statuses.

GUI element	Description
Groups	Add applications and departments to combine resources that are grouped by applications and departments. Applications and departments are used to monitor and view status, capacity, and performance information, and view details about their related resources. You can also complete the following actions to manage and administer applications and departments:
	Create applications and departments to model data for storage resources, allowing for enhanced capacity trending and performance troubleshooting.
	View the status of resources that make up the applications and departments.
	Create resource filters to automatically add resources to applications.
	Add resources directly to applications and add applications to existing applications and departments.
	Add departments to existing departments.
	View subcomponent (member application) information to support hierarchal levels of storage grouping. Paragraphs applications and departments.
	Remove applications and departments.
	Review the distribution of capacity in your storage environment to determine whether you have sufficient capacity to meet the demands of your tiered storage.
Advanced Analytics	Configure IBM Spectrum Control for reclaiming storage that is not being used.
	Discontinued support: Other advanced analytics features, such as cloud configuration, provisioning, and optimization, are no longer supported in IBM Spectrum Control. While these features might still work in this release, it's recommended that you use another tool for your provisioning and optimization needs, when possible. For a complete list of discontinued features, see "Discontinued features in IBM Spectrum Control" on page xxx.

GUI element	Description
Reports	Use the Reports option to create, configure, schedule, and send chargeback reports for applications, departments, hypervisors, and servers.
	Use IBM Spectrum Control when you want to individually view details about resources, for example when you want to troubleshoot a resource.
	Use the IBM Cognos Analytics reporting tool to analyze multiple storage systems, switches, servers, and hypervisors. For example, you can compare performance metrics across multiple storage systems. You can also view a report about the aggregate volume load that is on a server or hypervisor.
	Use the IBM Cognos Analytics reporting tool to schedule reports to run on different systems and with different options. You can specify schedules and output formats for the reports, and different ways to share the reports.
	For information about how to get started with the IBM
	Cognos Analytics reporting tool, see http://www.ibm.com/support/knowledgecenter/en/SSEP7J_11.0.0.
Settings	Complete the following actions to customize IBM Spectrum Control for your environment:
	Configure IBM Spectrum Control to send alert notifications by email, using SNMP traps, or using a Tivoli® Netcool®/OMNIbus server. An alert notification is sent when an alert condition is detected on the monitored resources in your environment.
	Create and modify alert policies to manage the alert definitions and notification settings that apply to different sets of resources.
	Modify the authentication repository and assign IBM Spectrum Control roles to user groups. Roles determine the product functions that are available to users.
	Specify how long to retain the data that is collected about resources and the log files that are generated by IBM Spectrum Control.
	Add and configure rollup servers so you can view capacity and status information about resources that are managed by your IBM Spectrum Control servers.

Keyboard navigation

Most of the features of the IBM Spectrum Control GUI are accessible by using the keyboard. For those features that are not accessible, equivalent function is available by using the command-line interface (CLI), except as noted in the product release notes.

You can use keys or key combinations to perform operations and initiate many menu actions that can also be done through mouse actions. The following sections describe the keys or key combinations for different parts of the GUI:

For navigating in the GUI and the context-sensitive help system:

- To navigate to the next link, button, or topic within a panel, press Tab.
- To move to the previous link, button, or topic within a panel, press Shift+Tab.
- To select an object, when the object is in focus, press Enter.

For actions menus:

- To navigate to the grid header, press Tab.
- To reach the drop-down field, press the Left Arrow or Right Arrow key.
- To open the drop-down menu, press Enter.
- To select the menu items, press the Up Arrow or Down Arrow key.
- To start the action, press Enter.

For filters:

To specify a filter option and text:

- 1. Press Tab to navigate to the magnifying glass icon.
- 2. Press the Up Arrow or Down Arrow key to navigate the filtering list.
- 3. Press Enter to select a filtering option.
- 4. When a filtering option is selected, the cursor moves to the filter text box. Type the filter text and press Enter. To reset a filter, press Enter.

For text fields:

- To navigate to text fields, press Tab.
- To navigate to the fields that are available for editing, press Tab.
- To navigate to the next field or to the **Submit** button, press Tab.

For tables or lists:

- To navigate between column headers, focus on a column header and use the Left Arrow and Right Arrow keys to move to other column headers.
- To navigate between data cells, focus on a data cell and use the Left, Right, Up, Down, Pageup, and Pagedown Arrow keys.
- To sort a column, focus on a column header and press Enter. The focus remains on the column header after the sort occurs.
- To change the size of a column, focus on the column header, hold Shift+Control, and press the Left or Right Arrow keys.
- To follow a link in a data cell, focus on a data cell and press Shift+F9.
- To open a menu for a table row, focus on the row and press Shift+F10.
- To select consecutive rows, select the first row and hold Shift, press the Up or Down Arrow keys to go to the last row in the range, and press the Space bar to add the new rows to the selection.
- To select non-consecutive rows, select a row and hold Control, press the Up or Down Arrow keys, and press the Space bar to add the new row to the selection.

Restriction: For Chinese languages, the keyboard combination Control+Space bar is not enabled for selecting multiple rows at the same time.

Keyboard navigation with Firefox for Mac users: If you're using Firefox on a Mac with IBM Spectrum Control and want to use keyboard navigation, complete the following steps:

- 1. In Firefox, go to Preferences > Advanced > General and clear the check mark for Always use the cursor keys to navigate within pages. This step enables the use of Tab key to navigate between GUI elements.
- 2. In the URL address bar of Firefox, type about:config and press Enter.
 - **Tip:** If a warning prompt is displayed, click the button to accept the risk of changing browser settings. Existing settings won't be changed; instead, you'll be adding a preference setting for accessibility.
- 3. To add an accessibility preference for tab focus, right-click on the configuration page and select New > Integer.
- 4. In the New integer value window, type accessibility.tabfocus and click OK.
- 5. Type 7 to set the integer value and click OK.
- 6. Open your Mac's System Preferences app, go to **Keyboard** > **Shortcuts**, and select **All Controls**.

Chapter 2. Social media for IBM Spectrum Control

Watch videos, read blogs, and explore social media to learn more about how to use IBM Spectrum Control to manage your storage environment. Be social! Join the conversation at #IBMStorage and #softwaredefinedstorage.

Videos for IBM Spectrum Control

Watch videos about how to use IBM Spectrum Control to manage your storage environment.

- "IBM Spectrum Control videos" on page 17
- "Cognos Analytics videos" on page 19

Important: These videos were recorded by using multiple versions of IBM Spectrum Control (formerly known as Tivoli Storage Productivity Center) and Cognos Analytics, but are at least partially applicable to the current version of the product.

IBM Spectrum Control videos

Important: The following videos were recorded by using multiple versions of IBM Spectrum Control (formerly known as Tivoli Storage Productivity Center), but are at least partially applicable to the current version of the product.

Simplifying storage, deploying new applications, and controlling costs with IBM Spectrum Storage See how IBM Spectrum Storage Suite can help you cope with data explosion, manage all types of storage, and quickly test and implement new cognitive applications, all while you maintain predictable, controlled costs.

(5:08) Simplify storage, deploy new applications, and control costs with IBM Spectrum Storage

IBM Storage / Spectrum Control overview

Watch a brief overview of IBM Spectrum Control and learn how it can help you manage a storage environment.

(1:08) FIBM Storage / Spectrum Control

IBM Spectrum Control from a technical point of view

Watch a brief overview of IBM Spectrum Control and learn about some of the latest features from a technical point of view.

(4:50) IBM Spectrum Control: Technical Point of View

Editing resource locations in bulk and use Cognos BI to report on capacity by location

See how to edit the locations of multiple resources at the same time and generate a "capacity by location" report in Cognos Analytics.

(2:05) Edit resource locations in bulk and use Cognos BI to report on capacity by location

Using chargeback reporting and reporting on storage consumption

Learn how to use IBM Spectrum Control for chargeback reporting and reporting on storage consumption; it will show how to set up items related to reports, for example, tiers and applications.

(7:20) How to use IBM Spectrum Control for chargeback reporting and reporting on storage consumption

Monitoring IBM Cloud Object Storage as an external storage tier for IBM Spectrum Scale

See how IBM Spectrum Control makes it easy to monitor the cloud storage that is used by your IBM Spectrum Scale clusters.

(2:54) Monitoring IBM Cloud Object Storage as an external storage tier for IBM Spectrum Scale

Determining when the capacity of a storage pool runs out

See how you can use the Zero Capacity metric to identify, based on space usage trends, when the capacity of a pool runs out.

Troubleshooting techniques with performance metrics

Learn how to use IBM Spectrum Control to troubleshoot an application and how to use the controls on the performance charts. Also, learn how to compare metrics across time ranges and identify the best or worst values for a performance metric.

(6:49) Learn how to use IBM Spectrum Control to troubleshoot performance problems in your storage environment

Gathering performance data into a performance support package

Learn how to quickly gather performance data in IBM Spectrum Control by using the Create Performance Support Package option.

(1:57) Gathering performance data in IBM Spectrum Control

Monitoring block storage performance

Learn how IBM Spectrum Control can be used to monitor the performance of block storage across your entire environment.

(6:08) Monitoring block storage performance with IBM Spectrum Control

Troubleshooting storage infrastructure performance

Learn how to monitor, compare, and troubleshoot performance on storage systems.

(10:30) Troubleshooting Storage Infrastructure Performance

Viewing and extracting performance data into a CSV file

Learn how to extract performance data from IBM Spectrum Control using the GUI.

(10:48) How to extract performance data

Monitoring IBM Spectrum Scale

Learn how to use IBM Spectrum Control to monitor multiple IBM Spectrum Scale clusters, view cross cluster relationships, monitor Transparent Cloud Tiering between IBM Spectrum Scale and IBM Cloud Object Storage, and view relationships between Network Shared Disks and SAN attached storage.

(5:33) How to monitor IBM Spectrum Scale with IBM Spectrum Control, V5.2.12

Monitoring the performance of IBM Spectrum Scale

Learn how to use IBM Spectrum Control to compare the performance of multiple IBM Spectrum Scale clusters and then analyze the clusters that have the heaviest I/O loads or the longest response times.

(5:38) Monitoring IBM Spectrum Scale performance with IBM Spectrum Control

Monitoring IBM Spectrum Scale clusters

Learn how to monitor IBM Spectrum Scale clusters.

(6:55) Using IBM Spectrum Control to monitor GPFS clusters

Alerting capabilities in IBM Spectrum Control

Learn about the alerting capabilities in IBM Spectrum Control.

(6:34) Alerting capabilities in IBM Spectrum Control

Cognos Analytics videos

Table 5. Videos about IBM Spectrum Control and Cognos Analytics (English only)		
Video	Summary	
(2:05) Migrating IBM Spectrum Control custom reports to IBM Cognos Analytics	See how to migrate IBM Spectrum Control custom reports from Cognos Business Intelligence to Cognos Analytics.	
(2:05) Cognos BI: Creating a capacity by location report in IBM Spectrum Control V5.2.15	Learn how to edit the locations of multiple resources at the same time and use Cognos Analytics to create a custom report which shows the storage capacity by location.	
(8:29) Cognos BI: Creating a tier capacity by location report	Learn how to create a capacity report of different tiers of storage across multiple data centers.	
(8:00) Cognos BI: Pre-defined reports	Learn how to use IBM Spectrum Control and Cognos BI to access pre-defined reports. These reports can help meet your SAN reporting needs in the areas of capacity, performance, and storage relationships.	
(15:54) Cognos BI: Setting up basic security for your reporting environment	Learn how to use the security features of Jazz for Service Management, Tivoli Common Reporting, and Cognos Analytics to provide some basic security for your reporting environment.	
(4:39) Cognos BI: Restricting user access to specific reports	Learn how to limit the reports and reporting capabilities that users can access within Cognos Analytics.	
(8:03) Cognos BI: Restricting which users can author and modify the directory structure that contains reports	Learn how to restrict which users can author and modify the directory structure that contains IBM Spectrum Control reports.	

Blogs for IBM Spectrum Control

Read blogs and follow us on Twitter to learn from the experts and get the latest tips about how to use IBM Spectrum Control to manage your storage environment.

IBM Spectrum Control has a strong presence on social media. Explore blogs and forums to read the perspectives of storage management experts and gain insights about how to use IBM Spectrum Control.

- Technical community
- Twitter:
 - #IBMStorage
 - #ibmsystems
 - #softwaredefined
 - Product news and announcements
 - Documentation hints and tips

Table 6. Recent blogs for IBM Spectrum Control	
Blog	Summary
IBM Spectrum Control™: Performance monitoring hints and tips	Get some hints and tips to help you use IBM Spectrum Control to monitor the performance of your storage.
IBM Spectrum Control: Performance monitoring for FlashSystem 840 / 900 don't forget the SNMP agent	Learn how to fix a performance monitoring problem for IBM FlashSystem 840 and IBM FlashSystem 900.
Whether Pure Flash Array monitoring or more capacity enhancements IBM Spectrum Control adds to your overall business uptime!	Learn about the features that were introduced in IBM Spectrum Control V5.3.7.
More, more and even more storage system monitoring and enhanced capacity reporting in IBM Spectrum Control Version 5.3.6!	Learn about the features that were introduced in IBM Spectrum Control V5.3.6.
More third-party storage support & IBM Spectrum Virtualize family enhancements highlight IBM Spectrum Control V5.3.5 release	Learn about the features that were introduced in IBM Spectrum Control V5.3.5.
IBM Storage Insights for IBM Spectrum Control IBM Cloud service highlights IBM Spectrum Control, V5.3.3!	Learn about the features that were introduced in IBM Spectrum Control V5.3.3.
Configuring alerts with alert policies highlights new release of IBM Spectrum Control!	Learn about the features that were introduced in IBM Spectrum Control V5.3.2.
New release of IBM Spectrum Control is now available: More reporting enhancements!	Learn about the features that were introduced in IBM Spectrum Control V5.3.1.
Easy report creation and IBM FlashSystem 9100 monitoring highlight new release of IBM Spectrum Control	Learn about the features that were introduced in IBM Spectrum Control V5.3.0.
New version of IBM Spectrum Control adds more monitoring information for IBM Spectrum Virtualize storage systems	Learn how IBM Spectrum Control provides more monitoring information for IBM Spectrum Virtualize, information on migrating reports to IBM Cognos Analytics, and integration with external LDAP repositories, among other features and enhancements.
IBM Spectrum Control virtual mailbox: Refresh my memory, where does IBM Spectrum Control keep log files for its main components?	Learn about the locations of the log files for IBM Spectrum Control. You can view these log files to identify where problems occur, and you can provide them to IBM Support so that they can help resolve any problems.
IBM Spectrum Control virtual mailbox: Is there a way to see capacity summaries for my storage systems?	Learn how to use the overview charts for storage systems to answer key physical capacity, volume capacity, and capacity savings questions.

Table 7. Blogs about IBM Spectrum Control V5.2.15	
Blog	Summary
Monitoring IBM Cloud Object Storage as an external storage tier for IBM Spectrum Scale	Learn how IBM Spectrum Control makes it easy to monitor the cloud storage that is used by your IBM Spectrum Scale clusters.
Edit resource locations in bulk and use Cognos BI to report on capacity by location	Learn how to edit the locations of multiple resources at the same time and use Cognos BI to report on capacity by location.
A great time-to-value feature: IBM Spectrum Control auto-creates your agentless servers	Learn how IBM Spectrum Control now automatically generates agentless servers for the physical servers and virtual machines in your SAN environment, so that you can easily see how they are consuming storage.
Agentless servers : how to handle an inconsistency in your storage environment	Learn how to merge agentless servers that represent the same physical server so that a single agentless server is shown in IBM Spectrum Control.
Pearls in the documentation	Delve into the IBM Spectrum Control Knowledge Center to discover information and tips that you might not be aware of but could find useful.

Table 8. Blogs about IBM Spectrum Control V5.2.14	
Blog	Summary
IBM Spectrum Control V5.2.14	Learn about the features that were introduced in IBM Spectrum Control V5.2.14.
Monitoring a stretched cluster or Hyperswap environment	Learn how to monitor your stretched cluster or Hyperswap environment with IBM Spectrum Control.
How does today's response time compare with last week's?	Learn how to compare performance metrics across different time ranges to help troubleshoot potential bottlenecks in an application.
Pearls in the documentation	Delve into the IBM Spectrum Control Knowledge Center to discover information and tips that you might not be aware of but could find useful.
IBM Spectrum Control virtual mailbox	Learn how to upload logs automatically and send them to IBM Support for investigation.

Table 9. Blogs about IBM Spectrum Control V5.2.13	
Blog	Summary
IBM Spectrum Control V5.2.13	Learn about the features that were introduced in IBM Spectrum Control V5.2.13.

Table 9. Blogs about IBM Spectrum Control V5.2.13 (continued)	
Blog	Summary
Chargeback and Consumer Reports: Don't show me the money!	Learn how to create a custom report that sends you a daily summary of the capacity used by applications without including storage costs.
Pearls in the documentation	Delve into the IBM Spectrum Control Knowledge Center to discover information and tips that you might not be aware of but could find useful.
IBM Spectrum Control virtual mailbox	Learn how much capacity is consumed by replication. As a storage administrator you are asked to provide a capacity report across your entire storage environment to your CTO. A critical part of capacity reporting is understanding how much capacity is consumed by replication.

Table 10. Blogs about IBM Spectrum Control V5.2.12	
Blog	Summary
IBM Spectrum Control V5.2.12	Learn about the features that were introduced in IBM Spectrum Control V5.2.12.
Consumer Reports: For your eyes only!	Learn how to configure a report that shows the block capacity and cost of the block capacity for a single storage consumer. Each consumer gets to star in their own report.
Chargeback and Consumer Reports: Same price for the same storage	In chargeback and consumer reports, you want to ensure that you charge the same price for the same type of storage whether it is block storage, copy data, or file storage. And, if your block storage is tiered, you want to charge the same price for each tier to all of your storage consumers.
Chargeback and Consumer Reports: How much storage was really eaten?	When you create chargeback and consumer reports, you can charge for the block storage that is allocated to the storage consumer or the block storage that is assigned to the storage consumer.
Look at that IBM FlashSystem data reduction!	FlashSystem A9000 and FlashSystem A9000R can reduce the size of your data significantly by using technologies such as compression and data deduplication. Learn how to see exactly what sort of savings you are getting, and how that compares to savings on other storage systems.
Pearls in the documentation	Delve into the IBM Spectrum Control Knowledge Center to discover information and tips that you might not be aware of but could find useful.

Table 10. Blogs about IBM Spectrum Control V5.2.12 (continued)	
Blog	Summary
From the IBM Spectrum Control virtual mailbox: Can I get enhanced optimization results for storage tiering and pool balancing?	You can reduce storage costs by moving volumes with low workloads from higher and more expensive tiers to lower and less expensive tiers. Storage performance can be improved and used more efficiently by moving volumes with heavy workloads to the tiers that meet their workload requirements. If you implement the recommendations, you can ensure that volumes are placed on the tiers that best match the workload requirements of the volumes.

Table 11. Blogs about IBM Spectrum Control V5.2.11	
Blog	Summary
IBM Spectrum Control V5.2.11	Learn about the features that were introduced in IBM Spectrum Control V5.2.11.
Chargeback Reports: Who's eating my storage?	Whether you manage or own applications, departments, hypervisors, or physical servers, you want to know how much storage is being consumed and how much it costs to maintain the storage that is consumed. And, you want that information in your inbox now, every week or month so that you can keep a close eye on storage costs and capacity consumption.
Capacity and compliance: Know your software license in IBM Spectrum Control	Your business organization just purchased a software license(s) for IBM Spectrum Control or any IBM product that includes IBM Spectrum Control such as IBM® Virtual Storage Center or IBM Spectrum Storage Suite and you want to assess your storage license situation.
Show your boss how your IBM FlashSystem compares	Learn how to compare the performance of your IBM FlashSystem to other storage systems. You can compare the response times and the I/O rates of IBM FlashSystem storage systems to other storage systems.
See how Cognos BI can help you create capacity and performance reports	Learn how to use the Cognos® Business Intelligence (BI) reporting tool to report on the condition, capacity, and performance of your storage resources.
From the IBM Spectrum Control hints and tips virtual mailbox: Is there a way to simplify the process of scheduling probes?	The automated probe option enables the creation of a probe and collects status and asset information more efficiently about your storage resources for monitoring.
Pearls in the documentation	Delve into the IBM Spectrum Control Knowledge Center to discover information and tips that you might not be aware of but could find useful.

Table 11. Blogs about IBM Spectrum Control V5.2.11 (continued)	
Blog	Summary
Navigating the new knowledge center	Learn how to quickly find information in the new IBM Knowledge Center and read about tips for getting the most out of its new interface.

Table 12. Blogs about IBM Spectrum Control V5.2.10	
Blog	Summary
IBM Spectrum Control V5.2.10	Learn about the features that were introduced in IBM Spectrum Control V5.2.10.
Extracting bulk performance data for analysis was never easier!	Learn about how to export bulk performance data for a storage system or fabric to a compressed file. If you contact IBM Support to help you analyze a performance problem with a resource, you might be asked to send this file.
IBM Spectrum Control: Rollup reporting provides a dynamic view into a large enterprise environment	Learn about how to use rollup reporting to gain a network-wide perspective of storage usage in your environment when you have multiple IBM Spectrum Control servers deployed.

Table 13. Blogs about IBM Spectrum Control V5.2.9	
Blog	Summary
IBM Spectrum Control V5.2.9	Learn about the features that were introduced in IBM Spectrum Control V5.2.9.

Table 14. Blogs about IBM Spectrum Control V5.2.8	
Blog	Summary
IBM Spectrum Control V5.2.8	Learn about the features that were introduced in IBM Spectrum Control V5.2.8.
Monitoring the entire data path of object storage systems	Learn about how to monitor object storage implementation on IBM Spectrum Scale and troubleshoot along the entire data path.

Chapter 3. Product licenses

Learn about the product licenses that are available for IBM Spectrum Control. Current licenses are IBM Spectrum Control and IBM Spectrum Control Select Edition. The IBM Spectrum Control Advanced Edition license is still applied, if you owned it in a previous release.

IBM Spectrum Control 5.4.0 offers two licenses based on the type of metric that you use:

- The IBM Spectrum Control license is based on the total number of storage capacity units (SCUs) that are required so that all of your capacity is licensed for all of your storage systems. This license applies if you currently have the IBM Spectrum Control or IBM Virtual Storage Center license.
- The IBM Spectrum Control Select Edition license (and the IBM Virtual Storage Center for Storwize license) is based on the total number of licenses that you need for your storage systems that use enclosures.

Note: Only one difference exists between the IBM Spectrum Control and IBM Spectrum Control Select Edition licenses: the license metric (SCUs versus enclosures). The features are otherwise identical in both solutions. To view the features, see Chapter 4, "Product overview," on page 41.

To view a complete list of devices that can be used with IBM Spectrum Control, go to http://www.ibm.com/support/docview.wss?uid=swg21386446.

IBM Storage Insights for IBM Spectrum Control: IBM Storage Insights for IBM Spectrum Control is an IBM Cloud service that can help you predict and prevent storage problems before they impact your business. It is complementary to IBM Spectrum Control and is available at no additional cost if you have an active license with a current subscription and support agreement for IBM Virtual Storage Center, IBM Spectrum Storage Suite, or any edition of IBM Spectrum Control.

For more information and to view a list of features in the service, see <u>IBM Storage Insights for IBM</u> Spectrum Control.

Related reference

Actions that are available based on role

Your IBM Spectrum Control role and product license determine the actions that are available in the product.

Definitions of licensing metrics

Learn more about the licensing metrics that are available for IBM Spectrum Control.

Three different license metrics are available for IBM Spectrum Control:

SCUs

Licenses are priced by the number of storage capacity units (SCUs) that you need for the capacity to be licensed for the storage system. The number of SCUs is based on the capacity and categories of drives that are used by the storage system. This metric is available for IBM Spectrum Control and IBM Virtual Storage Center.

Total TiB

Licenses are priced by the total capacity in tebibytes (TiB) for all storage systems that use the capacity license model. This metric is available for the IBM Virtual Storage Center Entry Edition and IBM Spectrum Storage Suite license.

Enclosures

Licenses are priced by the number of drive slots and capacity in the storage enclosures. This metric is available for IBM Spectrum Control Select Edition and IBM Virtual Storage Center for Storwize license.

For each enclosure, the number of licenses is calculated by using the following rules. Calculations are rounded up to the nearest whole number.

1. The number of slots in the enclosure divided by 25.

- 2. The sum of the physical capacity of the disks in the enclosure divided by 500 TiB.
- 3. The higher number from rule 1 and rule 2 is the number of licenses that are needed for the enclosure.

To learn about the licensing of IBM Spectrum Control, click the Announcement in <u>Table 15 on page 26</u>. Information about licenses is available in the section *Charge Metrics*.

Table 15. Product announcements for IBM Spectrum Control licenses					
License Announcement Revised Announcement					
IBM Spectrum Control - Storage Capacity Unit	5.3.0.0	5.3.6.0			
IBM Spectrum Control - Select/ Enclosure license	5.3.0.0	5.3.7.0			

Names of equivalent licenses in previous releases

IBM Spectrum Control licenses were known by different names in previous releases. Use this table to determine the names.

Table 16. IBM Spectrum Control licenses and their previous names			
Names of IBM Spectrum Control 5.4.0 licenses Names of IBM Spectrum Control 5.3 licenses			
IBM Spectrum Control IBM Spectrum Control Standard Edition			
IBM Spectrum Control Select Edition IBM Spectrum Control Standard Select Edition			

Actions that are available based on role

Your IBM Spectrum Control role and product license determine the actions that are available in the product.

Users who are assigned the Administrator role or the Monitor role can use product functions. The actions that are available for each function depend on the role that is assigned to the user:

Administrator role

Users who are assigned the Administrator role have access to all monitoring and administrative actions.

Monitor role

Users who are assigned the Monitor role can view information about monitored resources and other objects such as tasks, alerts, and service classes. They can acknowledge alerts and resource statuses, open logs, and open management GUIs.

The following table outlines the actions that are available only for the Administrators role. All other actions are available to the Monitor and Administrator roles. In addition to the restrictions listed in this table, users who are assigned the Monitor role do not have access to user management functions.

Table 17. Product ac	tions that are available only to users with the Administrator role
Function	Actions that require the Administrator role
Single dashboard view of the storage environment that you can use to manage storage systems, hypervisors, servers, and Fibre Channel fabrics.	 Adding and removing resources Administering connections Scheduling data collection Changing and viewing the automated probe schedule Viewing and editing history retention settings Modifying license settings
Performance monitoring for storage systems and Fibre Channel networks.	Scheduling performance monitors Starting or stopping performance monitors
Capacity and usage monitoring of resources.	 Scheduling probes Starting or stopping probes Modifying Storage Resource agents Enabling automatic zoning
Health and alerting for hypervisors, networks, servers, and storage systems.	 Creating, modifying, and deleting alert policies Setting which alert policy manages a resource Adding and modifying resources for management by an alert policy Defining and modifying alert definitions Editing alert notification settings
Capacity and performance of the storage that applications, departments, and general groups use.	 Creating applications, departments, and general groups Creating, modifying, and removing filters to add resources to applications Adding and removing resources in applications and general groups, directly Adding applications as members of other applications Adding departments to other departments Adding applications to departments
Storage reclamation	Viewing volumes that can be reclaimed
Roll-up reporting, in which capacity data is combined from multiple instances of IBM Spectrum Control for reporting purposes.	 Adding and removing subordinate servers Starting a probe for a subordinate server Modifying the connection information for a subordinate server

Table 17. Product ac	tions that are available only to users with the Administrator role (continued)
Function	Actions that require the Administrator role
Predefined Reports Predefined capacity reports allow users to quickly create reports about capacity anomalies and shortfalls, which can be scheduled and sent by email or saved to the user's file system, or both. Predefined inventory reports allow users to quickly create reports about their storage resources, which can be scheduled and sent by email or saved to the user's file system, or both.	Creating, deleting, and editing reports Configuring the email server Emailing reports Saving reports to the file system
Custom reports From any table view in the GUI, custom reports can be created, which can be scheduled and sent by email or saved to the user's file system, or both, about capacity of storage resources, the configuration and attributes of storage resources, and the performance of storage resources.	 Creating, deleting, and editing reports Configuring the email server Emailing reports Saving reports to the file system

Table 17. Product ac	tions that are available only to users with the Administrator role (continued)
Function	Actions that require the Administrator role
Chargeback and consumer reports Chargeback reports show the capacity and the cost of the storage that is used by applications, departments, hypervisors, and physical servers. Consumer reports show the capacity and the cost of the block storage that is used by an application, department, hypervisor, and physical server.	Creating, deleting, and editing reports Configuring the email server Emailing reports
Capacity limits for block storage systems and pools If your company has a policy to set a limit on the capacity that is used, you can set a capacity limit. When the capacity limit is set, you can then monitor the amount of capacity that is available before the capacity limit is reached.	Setting capacity limits Defining alerts for capacity limits Removing capacity limits

Related concepts

Product licenses

Learn about the product licenses that are available for IBM Spectrum Control. Current licenses are IBM Spectrum Control and IBM Spectrum Control Select Edition. The IBM Spectrum Control Advanced Edition license is still applied, if you owned it in a previous release.

Licensing

View the capacity of the storage systems that IBM Spectrum Control monitors in your environment. Then, check the records of the licenses that you purchased to determine how much additional storage can be monitored before you're no longer covered.

To determine how many licenses are needed for the capacity and enclosures that you monitor, complete these steps on the **Settings** > **Licensing** page:

- 1. Know your licenses.
- 2. Understand the licensing models.
- 3. Select the license model for each storage system.
- 4. <u>Update any settings that were not automatically updated or correctly identified for each storage system.</u>
- 5. Determine if you're covered.

1. Know your licenses

To get a list of the licenses that you purchased, complete one or more of these actions:

- Contact the person within your organization who originally purchased the licenses.
- Contact your IBM seller or IBM Business Partner.
- Go to Passport Advantage® Online.

2. Understand the licensing models

The licenses that you purchased determine which license models and values on the **Licensing** page are used to determine if you're covered. In some cases, a combination of values is used. For example, if you have some storage systems that are licensed by capacity and some that are licensed by enclosure, both the Total SCU and the Total Enclosure License values might be used.

License	License model
IBM Spectrum Control IBM Virtual Storage Center	Measured by the number of storage capacity units (SCUs) that are needed to cover the total usable capacity of the storage systems that IBM Spectrum Control monitors. For more information, see the description of the SCU value on the Licensing page.
IBM Spectrum Control Select Edition IBM Virtual Storage Center for Storwize	Measured by the number of licenses that are needed for the storage enclosures, modules, or expansions that IBM Spectrum Control monitors. See "Calculating the number of enclosure licenses for your storage systems" on page 34.
IBM Virtual Storage Center Entry Edition IBM Spectrum Storage Suite	Measured by the total capacity in tebibytes (TiB) that IBM Spectrum Control monitors.

3. Select the license model for each storage system

You must specify the license model for each of your storage systems. The license model, capacity or enclosure, is the type of IBM license that you purchased for the storage system.

License model	Description	
Capacity	The category of storage that is used by the storage system is needed to calculate the SCU usage. The category is based on the drive classes of the RAID array and is automatically assigned following a probe.	
	For IBM Spectrum Scale and virtual nodes in IBM Cloud Object Storage, you must specify the storage category.	
Enclosure	A probe automatically detects the number of enclosures that the storage system manages and determines the number of enclosure licenses that are required. You can manually update the number of enclosures licenses.	
	After you modify the number of enclosure licenses, that number will not change unless you manually update it again. Probes that detect a different number of enclosure licenses do not modify a number that was manually updated.	

Tip: For IBM System Storage DS8000, the license model is automatically assigned after a probe.


4. Update settings for each storage system

Update any settings that were not automatically updated or correctly identified for your storage systems. Examples of these settings include the number of enclosures in a storage system that were not identified or the drives in a storage system that were not categorized.

Learn more about configuring license settings: For more information about how to update the license settings for your storage systems, see "Configuring license values for storage systems" on page 37.

5. Determine if you're covered

To verify that your purchased licenses cover the capacity and enclosures that are monitored by IBM Spectrum Control, compare the values in your purchased licenses to the values at the top of the page.



- Look at **Total TiB** for these licenses: IBM Virtual Storage Center Entry Edition and IBM Spectrum Storage Suite.
- Look at **Total SCU** for these licenses: IBM Spectrum Control, IBM Virtual Storage Center, and older IBM Spectrum Control licenses that are based on SCU.
- Look at **Total Enclosure Licenses** for these licenses: IBM Spectrum Control Select Edition, IBM Virtual Storage Center for Storwize, and older IBM Spectrum Control licenses that are based on enclosure.

Learn more about license values: For more information about how capacity and enclosure values are calculated for storage systems and storage virtualization systems (such as IBM SAN Volume Controller), see "How licensing values are calculated" on page 34.

What's next

If the **Total TiB**, **Total SCU**, and **Total Enclosure Licenses** values are less than the value of your purchased licenses, you're covered.

If these values are more than the value of your licenses, you can reduce the amount of storage that you're monitoring or contact your IBM seller or IBM Business Partner to purchase additional storage.

Information about the summary values

The following, summary information is shown at the top of the **Licensing** page:

Capacity License Model

The number in parentheses shows the total number of storage systems that are licensed by storage capacity.

Total TiB

The sum of the capacity to be licensed for all storage systems that use the capacity license model. This value applies if you have the IBM Virtual Storage Center Entry Edition or IBM Spectrum Storage Suite license.

Total SCU

The total number of storage capacity units that you need for the capacity to be licensed for all storage systems that use the capacity license model. This value applies if you have the IBM Spectrum Control or IBM Virtual Storage Center license.

Category 1, 2, 3, or 4

The capacity breakdown of the storage environment by the categories of drives that are used by the storage systems. The breakdown is shown both as a bar chart and as tebibytes (TiB).

Category Info

For storage systems that are licensed by capacity rather than enclosure, a link that describes the categories of drives that are used by RAID arrays.

Enclosure License Model

The number in parentheses shows the total number of storage systems that are licensed by the enclosures that IBM Spectrum Control monitors.

Total Enclosure Licenses

The total number of licenses that are required for storage systems that use the enclosure license model. This value applies if you have the IBM Spectrum Control Select Edition or IBM Virtual Storage Center for Storwize license.

Information about storage system values

The following information is shown for each storage system on the **Licensing** page:

License Model

The type of IBM Spectrum Control license that was purchased for the storage system. Systems that use the capacity model are charged by the capacity to be licensed or the SCUs. Systems that use the enclosure model are charged by the number of enclosure licenses.

SCU

The number of storage capacity units that the license must cover for the storage system.

An SCU is the measure of capacity by which the license for a storage system is charged. The number of SCUs is based on the drive classes that are used by the storage system and the capacity that is assigned to each category. The following table shows the SCU categories, drive classes, and ratios that are supported by IBM Spectrum Control.

Table 18. Supported SCU licenses, drive classes, and ratios				
Category	Drive Classes	Ratio of usable storage to SCU		
1	Storage Class Memory (SCM) drives and managed disks on IBM Spectrum Virtualize for Public Cloud	1 TiB to 1 SCU		
2	Flash and Solid-State Drives (SSDs)	1 TiB to 0.847 SCU 1.18 TiB to 1 SCU		
3	10K Serial Attached SCSI (SAS) Drives, 15K Fibre Channel Drives, and storage systems that use Category 4 drives with advanced architectures	1 TiB to 0.5 SCU 2 TiB to 1 SCU		
4	Near Line SAS (NL-SAS) and Serial ATA (SATA) Drives	1 TiB to 0.25 SCU 4 TiB to 1 SCU		

Tips:

- · Any storage capacity that uses drive classes that are not listed in the table is classified as Category
- Calculations are rounded up to the nearest whole number.

The following table provides examples that you can use when you want to determine your SCU usage.

Table 19. Examples of SCU usage					
Storage category	TiB usable	SCU ratio	Total SCUs		
1	42	1	42		
2	400	1.18	339		
3	800	2	400°		
4	1600	4	400		
Totals	2842		1181		

Category

For capacity license models, the breakdown of the storage environment by the four categories of drives that are used by the storage system. The breakdown is shown as tebibytes (TiB).

For most storage systems, the storage category is set automatically. However, you must set the category for IBM Spectrum Scale and virtual appliances in IBM Cloud Object Storage. If the storage category is not set, IBM Spectrum Control assumes, for the calculations of SCU usage, that Category 1 drives are used.

For enclosure license models, the number of licenses that you need for the storage enclosures that the storage system manages. You can modify the number of enclosure licenses. For example, a probe might detect that two enclosure licenses are needed for a storage system, but you're aware of two external enclosures that were not detected. You calculate that these enclosures needs one license each. In that case, you can change the probe-detected value (2) to the more accurate number of enclosure licenses (4).

Keep in mind the following restrictions before you modify the number of enclosure licenses:

• You cannot reduce the number of enclosure licenses to be less than the number that was detected by a probe.

• After you modify the number of enclosure licenses, that number will not change unless you manually update it again. Probes that detect a different number of enclosure licenses do not modify a number that was manually updated.

Restriction:

- For SAN Volume Controller, if a probe doesn't detect any enclosures, **1** is displayed automatically.
- For XIV systems and IBM Spectrum Accelerate, the number of enclosures can't be edited.

Capacity to be Licensed (TiB)

The capacity that must be licensed for the storage system. Capacity that is reserved for formatting or RAID management is not included. For storage virtualizers and IBM Spectrum Scale, this value includes the capacity that is provided by back-end storage systems that are not being monitored by IBM Spectrum Control.

Capacity from back-end storage systems that are monitored by IBM Spectrum Control is recorded separately on the **Licensing** page and is not included in this value.

Monitored Back-end Capacity (TiB)

For storage virtualizers and IBM Spectrum Scale, the capacity that is provided by back-end storage systems that are being monitored by IBM Spectrum Control.

This capacity is not included in the capacity to be licensed of the virtualizer or IBM Spectrum Scale. Each monitored back-end storage system has a separate record on the **Licensing** page.

Unmonitored Back-end Capacity (TiB)

For storage virtualizers, the capacity that is provided by back-end storage systems that are not being monitored by IBM Spectrum Control. This capacity is included in the capacity to be licensed of the virtualizer. Unmonitored storage systems do not have their own record on the Licensing page.

Total Capacity (TiB)

The total capacity of the storage system, including all monitored and unmonitored back-end storage.

Total Capacity = Capacity to be Licensed + Monitored Back-end Capacity

Uncategorized Capacity (TiB)

The total capacity that is not categorized for the storage system. Capacity that is not categorized is also included in the Capacity to be Licensed value.

For most storage systems, IBM Spectrum Control automatically assigns the storage category for the pools on the storage system, based on the technology type of the drives in the pool. However, for IBM Spectrum Scale and virtual appliances on IBM Cloud Object Storage, IBM Spectrum Control cannot determine the storage category.

To ensure that the correct storage category is used, assign the storage category on the **Licensing** page. If the storage category is not set, IBM Spectrum Control assumes, for the calculations of SCU usage, that Category 1 drives are used.

Enclosure Licenses

The number of enclosure licenses that you need for the storage system.

How licensing values are calculated

Learn about how the capacity to be licensed and the number of enclosure licenses that are needed for your storage systems are calculated.

Calculating the number of enclosure licenses for your storage systems

For each enclosure, the number of licenses that are needed is calculated by using the following rules. Calculations are rounded up to the nearest whole number.

- 1. The number of slots in the enclosure divided by 25.
- 2. The sum of the physical capacity of the disks in the enclosure divided by 500 TiB.

3. The higher number from rule 1 and rule 2 is the number of licenses that are needed for the enclosure. Examples of the enclosure licensing calculations.

Rule 1 - Number of slots in the storage enclosures (or modules or expansions) that are managed:

- Enclosure with 12 slots requires 1 license.
- Enclosure with 92 slots requires 4 licenses.

Rule 2 - Sum of the physical capacity of the disks in the enclosure:

- Enclosure with 200 TiB requires 1 license.
- Enclosure with 1,250 TiB requires 3 licenses.

Example of a storage system with 4 enclosures:

Enclosure	Slots	Rule 1 - Licenses required	Capacity (TiB)	Rule 2 - Licenses required	Licenses required
Enclosure 1	28	2	400	1	2
Enclosure 2	92	4	3,700	8	8
Enclosure 3	60	3	1,750	4	4
Enclosure 4	16	1	225	1	1
				Total enclosure licenses required	15

Calculating the capacity to be licensed for your storage systems

Capacity to be licensed is the usable capacity of the storage systems that are being monitored and that is made available for storage consumption. Apart from IBM Spectrum Scale and IBM Cloud Object Storage, the capacity to be licensed is based on the usable capacity of the RAID arrays on the storage system.

Information about the calculation of capacity to be licensed:

- Capacity that is reserved for overheads, such as RAID management is not included.
- Data reductions savings at the storage system level, such as compression and deduplication, are not included. For storage systems that use data reduction technologies, the capacity to be licensed is the usable capacity, not the effective capacity.
- Capacity to be licensed is the total capacity of the storage system after RAID is applied and spare disks and capacity are deducted. So, capacity to be licensed is the capacity that can be provisioned to servers, when no over-provisioning is used. Any form of cache, for example, DRAM, SSD, or flash storage is not included, and the capacity to be licensed is less than the raw capacity of the storage system.

Refer to the following information for details about how capacity to be licensed is calculated for different storage systems:

Storage virtualization systems

Capacity to be licensed is calculated for the following storage virtualizers:

- Storage systems that run IBM Spectrum Virtualize, such as IBM SAN Volume Controller and IBM FlashSystem 9100.
- Non-IBM storage systems such as Hitachi VSP.

A storage virtualizer includes capacity from other storage systems. Capacity from back-end storage systems that are not monitored by IBM Spectrum Control is counted in the capacity to be licensed of the virtualizer. Capacity from back-end storage systems that are monitored by IBM Spectrum Control is not counted. The license for the monitored back-end storage is recorded separately and has its own record on the **Licensing** page.

For storage virtualizers that have internal capacity, such as IBM Storwize V7000, FlashSystem 7200, and IBM FlashSystem 9100, the capacity to be licensed is calculated as:

```
sum of the physical capacity of the internal RAID arrays + sum of the capacity of the external MDisks from unmonitored back-end systems
```

For storage virtualizers that do not have internal capacity, such as some models of IBM SAN Volume Controller, the capacity to be licensed is calculated as:

```
sum of the capacity of the external MDisks from unmonitored back-end systems
```

Block and file storage systems

The capacity to be licensed is the sum of the physical capacity of the RAID arrays on the storage system.

Exception: For FlashSystem A9000, FlashSystem A9000R, and IBM XIV Storage System, the capacity to be licensed is the overall physical capacity of the storage system.

For the following storage systems, the capacity and enclosure license information is not calculated because IBM Spectrum Control does not collect the RAID array information from the device. The value Unavailable is shown in the SCU and Capacity to be Licensed columns on the **Licensing** page.

- Block and file storage systems that are managed by SMI-S providers. For example:
 - DS4000, DS5000, and DS6000
 - NetApp Data ONTAP 8.1
 - Dell EMCVMAX, VNX, and VNXe
 - Hitachi VSP F series

To get a complete picture of your licensing requirements, complete the following tasks for these storage systems:

- If the storage system is licensed by enclosure, you must manually enter the number of enclosure licenses. The value is then automatically included in the Total Enclosure Licenses value in the summary area at the top of the page.
- If the storage system is licensed by capacity, you must manually calculate the SCUs that you need for the storage system. Then, add that value to the Total SCU value in the summary area.

IBM Spectrum Scale storage systems

The capacity to be licensed is the sum of the capacity of the Network Shared Disks (NSDs) that are assigned to the file system pools.

NSD capacity from back-end storage systems that are not monitored by IBM Spectrum Control is included.

The following capacity is not included in the capacity to be licensed:

- NSDs that are not assigned to a pool.
- NSD capacity from back-end storage systems that are monitored by IBM Spectrum Control. The license for the monitored back-end storage is recorded separately and has its own record on the **Licensing** page.
- NSD capacity on external pools.

For IBM Spectrum Scale storage systems that are configured for file and object storage, the capacity to be licensed for the object storage is included in the file system pool capacity and is not counted separately.

Don't forget to assign the storage category for your file system pools: The storage category is based on the technology type of the drives in the pool, for example, Storage Class Memory and Flash,

and is used in the calculation of SCU usage. To ensure that the correct storage category is used, assign the storage category on the **Licensing** page.

If the storage category is not set, IBM Spectrum Control assumes, for the calculations of SCU usage, that Category 1 drives are used.

Calculating enclosure licenses for IBM Spectrum Scale: For IBM Spectrum Scale storage systems that are licensed by enclosure, you need one enclosure license for each NSD server node that is connected to NSDs.

IBM Cloud Object Storage systems

The capacity to be licensed is the sum of the capacity of the COS Slicestor nodes in IBM Cloud Object Storage.

Don't forget to assign the storage category for your virtual COS Slicestor nodes: The storage category is based on the technology type of the drives in the node, for example, Storage Class Memory and Flash, and is used in the calculation of SCU usage. A COS Slicestor node can be a server or a virtual appliance. For the server nodes, IBM Spectrum Control automatically assigns the Near-Line drive category. For virtual appliances, IBM Spectrum Control cannot determine the storage category of the drives in the node.

For the calculation of SCU usage for the node, IBM Spectrum Control assumes that Category 1 drives are used in the virtual appliance. To ensure that the correct storage category is used in the calculations, assign the storage category on the **Licensing** page for each virtual node.

Calculating enclosure licenses for IBM Cloud Object Storage

To calculate the number of enclosure licenses, the same rules that are specified in <u>"Calculating</u> the number of enclosure licenses for your storage systems" on page 34 are used.

For the calculation:

- Each COS Slicestor node is treated as an enclosure.
- Instead of enclosure slots, the calculation uses the number of drives in the node.
- The capacity is based on the total capacity of the drives in the node.

Example of IBM Cloud Object Storage with 3 COS Slicestor nodes:

Nodes	Drives	Rule 1 - Licenses required	Capacity (TiB)	Rule 2 - Licenses required	Licenses required
Node 1	16	1	400	1	1
Node 2	64	3	2,400	5	5
Node 3	32	2	350	1	2
				Total enclosure licenses required	8

Configuring license values for storage systems

Discover the capacity that IBM Spectrum Control monitors so that you can determine whether you are in compliance with your license.

To determine the storage capacity that IBM Spectrum Control monitors, you must confirm the type of IBM license that you purchased for each storage system (known as the *license model*).

If you do not know the type of IBM Spectrum Control license that you purchased for each storage system, you can skip the process and return to complete the configuration later.

Whether you log on to IBM Spectrum Control for the first time or after an upgrade, the procedure for configuring license compliance is similar.

1. Click Settings > Licensing.

- If you are logging on to IBM Spectrum Control for the first time, the **Set License Model** page is displayed, with no selections made in the License Model column. Go to step "2" on page 38.
- If you are logging on to IBM Spectrum Control after an upgrade, the **Licensing** page is displayed, with the selections from the previous release in the License Model column. If some of your storage systems were not classified by license model and pool category in the last release, you must classify them by clicking **Assign licenses**.

Important: For some storage systems that support the IBM Easy TierIBM feature, the capacity is not calculated when you upgrade from IBM Spectrum Control. In that case, it is necessary to probe the storage systems again. Following an upgrade and before you open the **Licensing** page for the first time, be sure to run a probe of all of your monitored storage systems.

2. Choose the type of license model. For storage systems that are licensed by enclosure, enter the number of enclosure licenses that you need for the storage system.

Tip: A storage enclosure that is managed by IBM Spectrum Control is an independently powered, channel-attached device that stores data on magnetic disks or solid-state drives. An example is a disk controller and its expansion units. Each expansion unit is a separate enclosure. It can be the main controller that houses disk or solid-state drives, or the expansion chassis that houses extra disk or solid-state drives to expand the total capacity of the storage system.

- 3. The IBM Spectrum Control software automatically assigns the category of disk for the pools:
 - If all automatic assignments are correct, click Finish.
 - If one or more assignments are incorrect, click the correct category for each pool and click **Finish**.

When you finish entering information, the **Licensing** page is displayed. The page presents data about the storage in your environment. If the page shows uncategorized storage systems, you can click the **Assign licenses** link to enter its license model and drive category. The page also indicates the total tebibytes (TiB) of capacity that IBM Spectrum Control manages. Values are shown as tebibytes (TiB). 1 TiB is equal to 2⁴⁰ (1,099,511,627,776) bytes. To determine license compliance, compare the total TiB with your proof of entitlement that you received when you purchased your licenses.

Licensing examples

See how the capacity to be licensed and other values on the **Licensing** page are calculated for different storage systems and configurations.

In these examples, SVC-1 is an IBM SAN Volume Controller and is monitored by IBM Spectrum Control.

Example 1: Back-end capacity is not monitored

This example uses the following configuration:

- SVC-1 has no internal storage and does not receive capacity from storage systems other than the ones used in this example.
- Flash-1 is a Storwize V5000, which manages one enclosure with a drive capacity of 750 TiB. The drive capacity is configured as a 600 TiB RAID array.

Here's the situation:

- Flash-1 is not monitored by IBM Spectrum Control.
- Flash-1 provides 300 TiB of capacity to SVC-1.

The following values are shown on the **Licensing** page:

Storage System	Enclosure Licenses	,	Back-end	Unmonitored Back-end Capacity (TiB)	Total Capacity (TiB)
SVC-1	0	300	ln	300	300

Explanation:

The 300 TiB Flash-1 capacity that is provided to SVC-1 is shown as unmonitored back-end capacity and is included in the capacity to be licensed for SVC-1.

The remaining 300 TiB capacity on Flash-1 is not provided to storage systems that are monitored by IBM Spectrum Control and does not need to be licensed.

Example 2: Back-end capacity is monitored

Configuration:

- SVC-1 has no internal storage and does not receive capacity from storage systems other than the ones used in this example.
- Flash-1 is a Storwize V5000, which manages one enclosure with a drive capacity of 750 TiB. The drive capacity is configured as a 600 TiB RAID array.

Situation:

- Flash-1 is monitored by IBM Spectrum Control and is license by capacity.
- Flash-1 provides 300 TiB of capacity to SVC-1.
- SVC-1 receives an additional 800 TiB of capacity from other storage systems that are not monitored by IBM Spectrum Control.

The following values are shown on the **Licensing** page:

Storage System	Enclosure Licenses	Capacity to be Licensed (TiB)	Monitored Back-end Capacity (TiB)	Unmonitored Back-end Capacity (TiB)	Total Capacity (TiB)
Flash-1	0	600	0	0	600
SVC-1	0	800	300	800	1100

Example 3: Back-end capacity is monitored and the storage system is licensed by enclosure

Configuration:

- SVC-1 has no internal storage and does not receive capacity from storage systems other than the ones used in this example.
- Flash-1 is a Storwize V5000, which manages one enclosure with a drive capacity of 750 TiB. The drive capacity is configured as a 600 TiB RAID array.

Situation:

- Flash-1 is monitored by IBM Spectrum Control and is license by enclosure.
- Flash-1 provides 300 TiB of capacity to SVC-1.
- SVC-1 receives an additional 800 TiB of capacity from other storage systems that are not monitored by IBM Spectrum Control.

The following values are shown on the **Licensing** page:

Storage System	Enclosure Licenses	Capacity to be Licensed (TiB)	Monitored Back-end Capacity (TiB)	Unmonitored Back-end Capacity (TiB)	Total Capacity (TiB)
Flash-1	2	600	0	0	600
SVC-1	0	800	300	800	1100

Explanation:

- Because Flash-1 is licensed by enclosure, the capacity to be licensed is not included in the Total SCU and Total TiB values at the top of the **Licensing** page.
- To calculate the number of enclosure licenses, the Flash-1 drive capacity (750 TiB) is used, rather than the RAID array capacity (600 TiB).

Calculation of Enclosure Licenses for Flash-1

The higher number from rule 1 and rule 2 is used.

- Rule 1: 24 slots in the enclosure, which requires 1 enclosure license
- Rule 2: 750 TiB drive capacity, which requires 2 enclosure licenses (750 / 500, rounding up to the nearest whole number)

For more information, see <u>"Calculating the number of enclosure licenses for your storage systems" on page 34.</u>

Chapter 4. Product overview

IBM Spectrum Control provides a set tools for managing storage capacity, availability, alerts, events, performance, and resources. It can reduce the complexity of managing a storage environment by centralizing, simplifying, and optimizing storage tasks that are associated with storage systems, storage networks, performance troubleshooting, and capacity management.

IBM Spectrum Control can help you detect potential problems on your storage devices. For example, it can notify you when a server or storage system is running out of disk space or warn you of impending storage hardware failure so that you can prevent unnecessary system and application downtime.

IBM Spectrum Control provides the following benefits and functions:

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

Get a quick overview of the storage monitoring features that are available in IBM Spectrum Control:

Table 20. Overview of features in IBM Spectrum Control		
Features	Description	
Understanding your environment	A single dashboard view of your storage environment to help you manage storage devices.	
	Monitor IBM and non-IBM block storage, file storage, object storage, hypervisors, fabrics, switches.	
	Capacity monitoring of storage systems, pools, volumes, file systems, and file sets.	
	Performance monitoring of storage systems, Fibre Channel networks, applications, and departments.	
	Drill down performance workflows to troubleshoot bottlenecks.	
	Define groups of resources that represent applications or departments.	
	Explore storage and server virtualization relationships.	
	Explore replication relationships.	
Health and alerting	Show health status of hypervisors, networks, servers, and storage systems, and selected internal entities.	
	Alert on status changes and thresholds of capacity, performance, and error metrics.	
	Customizable, multi-conditional alerting and the ability to create alert policies for many devices.	

Features	Description	
Reporting	Inventory, capacity, and performance reports.	
	Reports that you can create from the information in the user-interface tables. You can configure, schedule, and save this information to your file systems. Or you can specify to send it to another user or person by mail.	
	Chargeback reports for showing the capacity and the cost of the storage that is used by applications, departments, hypervisors, and physical servers.	
	Consumer reports for showing the capacity and the cost of the block storage that is used by an application, department, hypervisor, and physical server.	
	Roll up reports in which capacity data is combined from multiple instances of IBM Spectrum Control for reporting purposes.	
	Reports that use REST API and Cognos.	
Analytics	For storage systems that run IBM Spectrum Virtualize, the transform storage feature moves volumes, and configures compression and the properties of the volumes.	
	Storage reclamation that you can realize when the block storage systems in your data center are analyzed, and you determine and reclaim volumes that aren't used for storing data. You can replace the existing storage space instead of purchasing new storage media.	
	Business impact analysis (applications, departments, and groups).	
General	Active directory and LDAP integration for managing users.	
	Customization options, including the ability to define the history retention of performance and capacity metadata, define how you are notified of alerts and notifications, and schedule when capacity and configuration metadata is automatically collected.	
	Inclusion of the license for IBM Copy Services Manager to manage 2-site replication, 3-site replication, and advanced copy services. Learn more about how to download and install Copy Services Manager at http://www.ibm.com/support/knowledgecenter/SSESK4 .	
	Snapshot-based protection that is provided by IBM Spectrum Protect™ Snapshot for the applications and databases that IBM Spectrum Control uses.	

Architecture

The IBM Spectrum Control consists of several components that form the infrastructure of its storage-management functions.

Data server

This component is the control point for product scheduling functions, configuration, and event information. 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.

Web server

IBM Spectrum Control uses WebSphere Application Server Liberty as the web application server to host its GUI, and storage management API for cloud.

This component manages the complex event processing that is related to alerting on the condition of resources and their attributes. The ability to detect and be notified about configuration, capacity, and performance changes within a storage environment is important to helping you maintain and administer storage resources.

Export server

This component enables the export of data from the tables in the GUI to HTML format and the creation of custom reports that show information about capacity, configuration, and health status of the resources in your storage environment. It runs in the Node, is environment and you can produce custom reports for storage systems, back-end storage systems, and their internal resources such as disks, pools, and volumes.

IBM Cognos Analytics (optional)

If you plan to use the optional Cognos Analytics reports, you must install IBM Cognos Analytics. You can use Cognos Analytics to create reports and view predefined reports to analyze multiple storage systems, switches, servers, and hypervisors. You can install Cognos Analytics on the same computer with IBM Spectrum Control or you can install Cognos Analytics on a separate computer.

Database

A single database instance serves as the repository for all IBM Spectrum Control components.

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

Use the IBM Spectrum Control GUI to manage the storage infrastructure in an enterprise environment.

CLI

Use the command-line interface (CLI) to issue commands for key IBM Spectrum Control functions.

Interfaces for IBM Spectrum Control

IBM Spectrum Control provides multiple user interfaces for managing the storage infrastructure in an enterprise environment.

IBM Spectrum Control GUI

This interface runs in a web browser and includes the ability to monitor, manage, and troubleshoot storage resources. You can access this interface from anywhere that you have a web browser and connectivity to a network.

IBM Cognos Analytics (optional)

This interface runs in a web browser. Use this interface to view predefined reports and create custom reports about the storage systems managed by IBM Spectrum Control. For more information, see https://www.ibm.com/support/knowledgecenter/en/SSEP7J 11.0.0.

Command-line interface

Use this interface to run IBM Spectrum Control commands from a command prompt. For more information about CLI commands, go to the product documentation at December Command-line interface.

Related concepts

"Managing resources" on page 111

Use IBM Spectrum Control to monitor and manage the resources and storage infrastructure in an enterprise environment. In the GUI, you can view the condition, capacity, performance, and relationships of storage resources.

Starting IBM Spectrum Control

You can start IBM Spectrum Control by opening a web browser and entering a web address for the IBM Spectrum Control logon page. For example, you might enter https://storage.example.com:9569/srm.

Start the IBM Spectrum Control GUI to administer and monitor the condition, capacity, and relationships of the resources within your storage environment.

1. On a server running the Windows operating system, start **IBM Spectrum Control GUI**. If you are not on a server running the Windows operating system, start a web browser and enter the following address in the address field:

```
https://host_name:port/srm
```

In the preceding address, specify the following values:

host name

The IBM Spectrum Control server. You can specify the host name as an IP address or a Domain Name System (DNS) name.

port

The port number for IBM Spectrum Control. The default port number for connecting to IBM Spectrum Control by using the HTTPS protocol is 9569. However, this port number might be different for your site. For example, the port number might be different if the default port range was not accepted during installation. If the default port number does not work, ask your IBM Spectrum Control administrator for the correct port number.

Tip: If you have a non-default port, check the value of the WC_defaulthost_secure property in installation_dir/web/conf/portdef.props file.

2. From the IBM Spectrum Control logon page, type your user name and password and click **Log in**. The GUI opens in the browser.

Tip: If you want to log on to the GUI with Windows Domain credentials, use this form: domain_name \user.

Navigation

IBM Spectrum Control provides many of the functions for managing a storage environment. To access the functions in its GUI, use the menu bar at the top of the main window.

GUI element	Description
Page banner	Use the banner at the top of every page to complete the following tasks:
	• View the name of the product license that is active.
	 Access resource pages and main product functions from the menu bar.
	 View a list of failed tasks by using the failed tasks icon You can retry or cancel those tasks.
	View a list of tasks that are running by using the running tasks icon tasks are running.
	 View the ID and role of the user who is logged in. The role of a user determines the product functions that are available to that user. For more information about roles, see "Role-based authorization" on page 80.
	 Access the online help to view information about the currently displayed page and the overall product.
	• Enter and exit rollup mode by using the rollup mode icon icon. You must add at least one subordinate rollup server for the rollup mode icon to display in the banner.

GUI element	Description
Home	Dashboard Use the dashboard to view the overall status of monitored resources and identify potential problem areas in a storage environment. You can learn about:
	 Condition and usage of resources
	 Entities that consume storage on those resources
	 Number and status of unacknowledged alert conditions that are detected on the monitored resources
	 Most active storage systems in your environment.
	Performance Monitors Use performance monitors to collect information about the performance of storage systems and switches. This information includes key performance metrics and the percentage of data collections that succeeded during the most recent 24 hours when the performance monitor was active.
	Alerts
	Use alerts to be notified when certain configuration, status, and performance conditions are detected on monitored resources.
	Tasks
	Manage the tasks that are created when you complete the following actions:
	 Place volumes on the tiers that match the workload requirements of the volume.
	 Move or convert volumes.
	 Move volumes from overutilized pools to pools that are less used.
	System Management View information about the overall condition of IBM Spectrum Control. You can view information about the servers on which the product is installed. This information includes component server and database status, certain server alerts and database connection alerts, server file-system capacity information, and remote volume-performance information.

GUI element	Description	
Storage	Add storage systems for monitoring and view asset, status, and performance information about their internal and related resources. You can also complete the following actions to manage and administer storage systems:	
	Modify the schedules for data collection.	
	 Place volumes on the tiers that match the workload requirements of the volume. 	
	Move or convert volumes.	
	 Move volumes from overutilized pools to pools that are less used. 	
	Add resources to a capacity pool.	
	Open the management GUIs for storage systems.	
	Update login credentials.	
	Acknowledge conditions and statuses.	
Servers	Add servers and hypervisors for monitoring and view asset and status information about their internal and related resources. You can also complete the following actions to manage and administer these resources:	
	 Deploy Storage Resource agents to servers for full monitoring. 	
	Modify the schedules for data collection.	
	Modify and upgrade Storage Resource agents.	
	Acknowledge conditions and statuses.	
Network	Add switches and fabrics for monitoring and view asset, status, and performance information about their internal and related resources. You can also complete the following actions to manage and administer these resources:	
	Modify the schedules for data collection.	
	Setting or modifying the zoning policy.	
	Open the management GUI for a switch.	
	Acknowledge conditions and statuses.	

GUI element	Description	
Groups	Add applications and departments to combine resources that are grouped by applications and departments. Applications and departments are used to monitor and view status, capacity, and performance information, and view details about their related resources. You can also complete the following actions to manage and administer applications and departments:	
	Create applications and departments to model data for storage resources, allowing for enhanced capacity trending and performance troubleshooting.	
	View the status of resources that make up the applications and departments.	
	Create resource filters to automatically add resources to applications.	
	Add resources directly to applications and add applications to existing applications and departments.	
	Add departments to existing departments.	
	View subcomponent (member application) information to support hierarchal levels of storage grouping. Paragraphs applications and departments.	
	Remove applications and departments.	
	Review the distribution of capacity in your storage environment to determine whether you have sufficient capacity to meet the demands of your tiered storage.	
Advanced Analytics	Configure IBM Spectrum Control for reclaiming storage that is not being used.	
	Discontinued support: Other advanced analytics features, such as cloud configuration, provisioning, and optimization, are no longer supported in IBM Spectrum Control. While these features might still work in this release, it's recommended that you use another tool for your provisioning and optimization needs, when possible. For a complete list of discontinued features, see "Discontinued features in IBM Spectrum Control" on page xxx.	

GUI element	Description
Reports	Use the Reports option to create, configure, schedule, and send chargeback reports for applications, departments, hypervisors, and servers.
	Use IBM Spectrum Control when you want to individually view details about resources, for example when you want to troubleshoot a resource.
	Use the IBM Cognos Analytics reporting tool to analyze multiple storage systems, switches, servers, and hypervisors. For example, you can compare performance metrics across multiple storage systems. You can also view a report about the aggregate volume load that is on a server or hypervisor.
	Use the IBM Cognos Analytics reporting tool to schedule reports to run on different systems and with different options. You can specify schedules and output formats for the reports, and different ways to share the reports.
	For information about how to get started with the IBM
	Cognos Analytics reporting tool, see http://www.ibm.com/support/knowledgecenter/en/SSEP7J_11.0.0.
Settings	Complete the following actions to customize IBM Spectrum Control for your environment:
	Configure IBM Spectrum Control to send alert notifications by email, using SNMP traps, or using a Tivoli® Netcool®/OMNIbus server. An alert notification is sent when an alert condition is detected on the monitored resources in your environment.
	Create and modify alert policies to manage the alert definitions and notification settings that apply to different sets of resources.
	Modify the authentication repository and assign IBM Spectrum Control roles to user groups. Roles determine the product functions that are available to users.
	Specify how long to retain the data that is collected about resources and the log files that are generated by IBM Spectrum Control.
	Add and configure rollup servers so you can view capacity and status information about resources that are managed by your IBM Spectrum Control servers.

Keyboard navigation

Most of the features of the IBM Spectrum Control GUI are accessible by using the keyboard. For those features that are not accessible, equivalent function is available by using the command-line interface (CLI), except as noted in the product release notes.

You can use keys or key combinations to perform operations and initiate many menu actions that can also be done through mouse actions. The following sections describe the keys or key combinations for different parts of the GUI:

For navigating in the GUI and the context-sensitive help system:

- To navigate to the next link, button, or topic within a panel, press Tab.
- To move to the previous link, button, or topic within a panel, press Shift+Tab.
- To select an object, when the object is in focus, press Enter.

For actions menus:

- To navigate to the grid header, press Tab.
- To reach the drop-down field, press the Left Arrow or Right Arrow key.
- To open the drop-down menu, press Enter.
- To select the menu items, press the Up Arrow or Down Arrow key.
- To start the action, press Enter.

For filters:

To specify a filter option and text:

- 1. Press Tab to navigate to the magnifying glass icon.
- 2. Press the Up Arrow or Down Arrow key to navigate the filtering list.
- 3. Press Enter to select a filtering option.
- 4. When a filtering option is selected, the cursor moves to the filter text box. Type the filter text and press Enter. To reset a filter, press Enter.

For text fields:

- To navigate to text fields, press Tab.
- To navigate to the fields that are available for editing, press Tab.
- To navigate to the next field or to the **Submit** button, press Tab.

For tables or lists:

- To navigate between column headers, focus on a column header and use the Left Arrow and Right Arrow keys to move to other column headers.
- To navigate between data cells, focus on a data cell and use the Left, Right, Up, Down, Pageup, and Pagedown Arrow keys.
- To sort a column, focus on a column header and press Enter. The focus remains on the column header after the sort occurs.
- To change the size of a column, focus on the column header, hold Shift+Control, and press the Left or Right Arrow keys.
- To follow a link in a data cell, focus on a data cell and press Shift+F9.
- To open a menu for a table row, focus on the row and press Shift+F10.
- To select consecutive rows, select the first row and hold Shift, press the Up or Down Arrow keys to go to the last row in the range, and press the Space bar to add the new rows to the selection.
- To select non-consecutive rows, select a row and hold Control, press the Up or Down Arrow keys, and press the Space bar to add the new row to the selection.

Restriction: For Chinese languages, the keyboard combination Control+Space bar is not enabled for selecting multiple rows at the same time.

Keyboard navigation with Firefox for Mac users: If you're using Firefox on a Mac with IBM Spectrum Control and want to use keyboard navigation, complete the following steps:

- 1. In Firefox, go to Preferences > Advanced > General and clear the check mark for Always use the cursor keys to navigate within pages. This step enables the use of Tab key to navigate between GUI elements.
- 2. In the URL address bar of Firefox, type **about:config** and press Enter.
 - **Tip:** If a warning prompt is displayed, click the button to accept the risk of changing browser settings. Existing settings won't be changed; instead, you'll be adding a preference setting for accessibility.
- 3. To add an accessibility preference for tab focus, right-click on the configuration page and select New > Integer.
- 4. In the New integer value window, type accessibility.tabfocus and click OK.
- 5. Type 7 to set the integer value and click **OK**.
- 6. Open your Mac's System Preferences app, go to **Keyboard** > **Shortcuts**, and select **All Controls**.

IBM Spectrum Storage Suite

IBM Spectrum Control provides monitoring, automation and analytics for multiple-vendor storage environments and is a member of IBM Spectrum Storage[™] Suite.

Licensing options

IBM Spectrum Storage Suite gives you unlimited access to the IBM Software Defined Storage product portfolio with licensing on a flat, cost-per-TB basis to make pricing easy to understand and predictable as storage capacity grows. You can save up to 40 percent compared with licensing the products separately. Non-production use of the software in test environments is included.

Learn more: Go to IBM Spectrum Storage Suite at IBM Marketplace.

Which products are included with IBM Spectrum Storage Suite?

Besides the aforementioned IBM Spectrum Control, the following products are included in the IBM Spectrum Storage Suite:

Table 21. Products that are members of IBM Spectrum Storage Suite			
Product	Key benefits	Links	
IBM® Spectrum Protect™	Protects your data with multi-site replication and flexible restore capacity, and reduces your backup costs up to 53 percent. It can simplify data protection where data is hosted in physical, virtual, software-defined or cloud environments.	Learn more and buy Product documentation	

Product	Key benefits	Links
IBM® Spectrum Protect™ Plus	Provides a data protection and availability solution for virtual environments that can be deployed in minutes and protect your environment within an hour. You can either implement it as a stand-alone solution, or integrate the solution with the IBM Spectrum Protect™ environment to offload copies for long-term storage and governance with scale and efficiency.	Learn more and buy Product documentation
	IBM Spectrum Protect Plus on IBM Cloud is a data protection and availability solution for virtual environments that can be deployed in minutes and protect your environment within an hour. IBM Spectrum Protect Plus on IBM Cloud can be provisioned with push-button simplicity to an existing or new VMware environment from the IBM Cloud console.	
IBM® Spectrum Scale [™]	Supports big data analytics and clustered applications with high-performance, scalable storage that enables global collaboration, simplifies workflows and lowers costs with cloud tiering.	Learn more and buy Product documentation
IBM® Spectrum Virtualize™	Provides an ideal way to manage and protect the huge volumes of data organizations use for bigdata analytics and new cognitive workloads.	Learn more and buy Product documentation
IBM Spectrum Virtualize for Public Cloud	Provides a way to transition existing IT architecture to a hybrid-cloud or cloud-based model, where servers, storage, and network infrastructure are delivered in a public cloud environment.	Learn more and buy Product documentation
IBM® Spectrum Archive™	Provides a direct, intuitive and graphical access to data stored in IBM tape drives and libraries by incorporating the Linear Tape File System™ (LTFS) format standard for reading, writing and exchanging descriptive metadata on formatted tape cartridges.	Learn more and buy Product documentation

Table 21. Products that are members of IBM Spectrum Storage Suite (continued)			
Product	Key benefits	Links	
IBM® Spectrum Accelerate™	Provides a software-defined block storage solution that is designed for rapid and flexible deployment across heterogeneous infrastructure on and off premises. It can be deployed on your choice of servers: x86 commodity-choice hardware, integrated appliances (including IBM XIV Storage System Gen3 and pre-validated third party appliances), and IBM public cloud.	Learn more and buy Product documentation	
IBM® Cloud Object Storage	Provides highly scalable software-defined on-premises storage system designed to store, protect and easily access unstructured data. Deploy in your own data centers using IBM appliances or IBM software running on industry-standard, certified hardware.	 Learn more and buy Product documentation 	
IBM Spectrum™ Discover	Modern metadata management software that provides data insight for petabyte-scale unstructured storage. The software easily connects to IBM Cloud Object Storage and IBM Spectrum Scale to rapidly ingest, consolidate and index metadata for billions of files and objects. It provides a rich metadata layer that enables storage administrators, data stewards and data scientists to efficiently manage, classify and gain insights from massive amounts of unstructured data.	Learn more and buy Product documentation	

You might also be interested in the following products

Table 22. Other products of interest					
Product	Key benefits	Links			
IBM Storage Insights Pro/Storage Insights	An IBM Cloud based product that provides capabilities to help optimize your infrastructure with features, such as capacity planning, reclamation and data tier planning, and performance troubleshooting at a low monthly subscription price.	Learn more and buy Product documentation			
	IBM Storage Insights, is a new application in the Storage Insights family. It is a no-cost support tool that is available to all organizations that have block storage systems.				
	The enhanced dashboard enables you to quickly monitor and assess the basic health, status, and performance of block storage systems. If a problem is detected, you can get help to investigate and troubleshoot the problem and minimize the impact of hardware and software issues before they impact the performance of your critical business applications.				
IBM Spectrum™ Copy Data Management	Makes copies available to data consumers when and where they need them, without creating unnecessary copies or leaving unused copies on valuable storage. It catalogs copy data from across your local and hybrid cloud and off-site cloud infrastructure, identifies duplicates, and compares copy requests to existing copies.	Learn more and buy Product documentation			

IBM Storage Insights for IBM Spectrum Control

IBM Storage Insights for IBM Spectrum Control is an IBM Cloud service that can help you predict and prevent storage problems before they impact your business. It is complementary to IBM Spectrum Control and is available at no additional cost if you have an active license with a current subscription and support agreement for IBM Virtual Storage Center, IBM Spectrum Storage Suite, or any edition of IBM Spectrum Control.

As an on-premises application, IBM Spectrum Control doesn't send the metadata about monitored devices offsite, which is ideal for dark shops and sites that don't want to open ports to the cloud. However, if your organization allows for communication between its network and the cloud, you can use

IBM Storage Insights for IBM Spectrum Control to transform your support experience for IBM block storage.

Why use IBM Storage Insights for IBM Spectrum Control

IBM Storage Insights for IBM Spectrum Control and IBM Spectrum Control work hand in hand to monitor your storage environment. Here's how IBM Storage Insights for IBM Spectrum Control can transform your monitoring and support experience:

- Open, update, and track IBM Support tickets easily for your IBM block storage devices.
- Get hassle-free log collection by allowing IBM Support to collect diagnostic packages for devices so you
 don't have to.
- Use Call Home to monitor devices, get best practice recommendations, and filter events to quickly isolate trouble spots.
- Leverage IBM Support's ability to view the current and historical performance of your storage systems and help reduce the time-to-resolution of problems.

To compare the features of **IBM Spectrum Control** and **IBM Storage Insights for IBM Spectrum Control**, check out the "Feature comparison" on page 57.

Important: To monitor the devices that you already monitor in IBM Spectrum Control, you must add them separately to IBM Storage Insights for IBM Spectrum Control. You must also deploy a data collector in IBM Storage Insights for IBM Spectrum Control to collect advanced metadata about devices.

Alerting tip: Alerts are a good way to be notified of conditions and potential problems that are detected on your storage. If you use **IBM Spectrum Control** and **IBM Storage Insights for IBM Spectrum Control** together to enhance your monitoring capabilities, it's recommended that you define alerts in one of the offerings and not both. By defining all your alerts in one offering, you can avoid receiving duplicate or conflicting notifications when alert conditions are detected.

How to get IBM Storage Insights for IBM Spectrum Control

To register for IBM Storage Insights for IBM Spectrum Control, go to https://www.ibm.com/it-infrastructure/storage/storage-insights/registration and follow the directions for current IBM Spectrum Control customers.

To see how easy it is to sign up and deploy a data collector, check out <u>Before you begin checklist for IBM</u> Storage Insights.

Questions and answers

Find answers to other questions about IBM Storage Insights for IBM Spectrum Control.

What's the difference between IBM Storage Insights, IBM Storage Insights Pro, and IBM Storage Insights for IBM Spectrum Control?

- **IBM Storage Insights** is an off-premises, IBM Cloud service that is available free of charge if you own IBM block storage systems. It provides a unified dashboard for IBM block storage systems with a diagnostic events feed, a streamlined support experience, and key capacity and performance information.
- IBM Storage Insights Pro is an off-premises, IBM Cloud service that is available on subscription and expands the capabilities of IBM Storage Insights. You can monitor IBM file, object, and software-defined storage (SDS) systems, and non-IBM block and file storage systems such as Dell EMC storage systems. It also includes configurable alerts and predictive analytics that help you to reduce costs, plan capacity, and detect and investigate performance issues. You get recommendations for reclaiming unused storage, recommendations for optimizing the placement of tiered data, capacity planning analytics, and performance troubleshooting tools.
- IBM Storage Insights for IBM Spectrum Control is similar to IBM Storage Insights Pro in capability and is available for no additional cost if you have an active license with a current subscription and

support agreement for IBM Virtual Storage Center, IBM Spectrum Storage Suite, or any edition of IBM Spectrum Control.

I already have a paid subscription to IBM Storage Insights Pro, does this service apply to me? Yes! When your next renewal date for IBM Storage Insights Pro approaches, contact your IBM Sales representative to see if a move to IBM Storage Insights for IBM Spectrum Control is right for you.

Tip: If you want to cancel your subscription to IBM Storage Insights Pro and switch to IBM Storage Insights for IBM Spectrum Control, check out the IBM Cloud Service Agreement at https://www.ibm.com/support/customer/csol/contractexplorer/cloud/csa/us-en/9 for cancellation information.

If I already have IBM Storage Insights (the free service), can I upgrade it to IBM Storage Insights for IBM Spectrum Control?

Yes, if you have an active license for IBM Spectrum Control, you can upgrade IBM Storage Insights to IBM Storage Insights for IBM Spectrum Control. Just go to https://www.ibm.com/it-infrastructure/storage-insights/registration, choose the option for IBM Spectrum Control, and follow the prompts.

Because IBM Storage Insights for IBM Spectrum Control is similar to IBM Storage Insights Pro in capability, does it use the same Service Level Agreement (SLA)?

No, IBM Storage Insights for IBM Spectrum Control doesn't include the service level agreement for IBM Storage Insights Pro. Terms and conditions for IBM Storage Insights for IBM Spectrum Control are available at http://www.ibm.com/software/sla/sladb.nsf/sla/sd-8410-01.

Do IBM Spectrum Control and IBM Storage Insights for IBM Spectrum Control share the metadata that they collect about storage?

No, metadata cannot be shared between the offerings or exported from one offering to the other.

Are the devices that I monitor in IBM Spectrum Control automatically monitored by IBM Storage Insights for IBM Spectrum Control?

To monitor the devices that you already monitor in IBM Spectrum Control, you must also add them to IBM Storage Insights for IBM Spectrum Control.

To collect advanced metadata about devices, you must also deploy a data collector in IBM Storage Insights for IBM Spectrum Control. Storage Resource agents and device connections that are defined in IBM Spectrum Control aren't used in IBM Storage Insights for IBM Spectrum Control.

To see which devices that you can monitor in IBM Storage Insights for IBM Spectrum Control, check out the list for IBM Storage Insights Pro in Supported resources.

If I use IBM Storage Insights Pro or IBM Storage Insights and switch to IBM Storage Insights for IBM Spectrum Control, is my existing metadata migrated?

Yes, the metadata that you collected in IBM Storage Insights Pro and IBM Storage Insights is automatically available in IBM Storage Insights for IBM Spectrum Control.

How is licensing handled?

IBM Storage Insights for IBM Spectrum Control uses the same licensing that you already have in place for IBM Spectrum Control. IBM Spectrum Control can be licensed by capacity, by storage capacity units (SCU), or by number of slots in enclosures.

How long can I use IBM Storage Insights for IBM Spectrum Control?

You can use IBM Storage Insights for IBM Spectrum Control for as long as you have an active license with a current subscription and support agreement for IBM Spectrum Control license. If your subscription and support lapses, you're no longer eligible for IBM Storage Insights for IBM Spectrum Control.

If your subscription and support lapses, don't worry. To continue using IBM Storage Insights for IBM Spectrum Control, simply renew your IBM Spectrum Control license.

You can also choose to subscribe to IBM Storage Insights Pro. For information about how to subscribe, see Want to try or buy IBM Storage Insights Pro?

Because IBM Storage Insights for IBM Spectrum Control is a cloud service, I'm concerned about security. Where can I find out more information about its security measures?

IBM Storage Insights for IBM Spectrum Control runs in IBM Cloud and adheres to IBM's rigorous security standards. For more information, check out the IBM Storage Insights Security Guide.

Where can I learn more about IBM Storage Insights?

Explore these links for more detailed information about IBM Storage Insights:

- IBM Storage Insights Knowledge Center
- IBM Storage Insights FAQ

How do I get support for IBM Storage Insights for IBM Spectrum Control?

To contact IBM Support for help and report issues that you encounter in IBM Storage Insights for IBM Spectrum Control, follow these steps:

- 1. Open a support case against IBM Storage Insights.
- 2. Describe the problem. To help us troubleshoot, include the URL of your IBM Storage Insights for IBM Spectrum Control instance.
- 3. Submit the case.

Feature comparison

Compare the features of IBM Spectrum Control and IBM Storage Insights for IBM Spectrum Control.

Table 23. Features in IBM Spectrum Control and IBM Storage Insights for IBM Spectrum Control				
Resource Management	Features	IBM Spectrum Control (Advanced edition)	IBM Storage Insights for IBM Spectrum Control	
Monitoring	Inventory	IBM and non-IBM block storage, file storage, object storage, hypervisors, fabrics, switches	IBM and non-IBM block storage, file storage, and object storage	
	Call Home events		~	
	Performance	~	~	
		(1-minute intervals)	(5-minute intervals)	
	Capacity	~	~	
	Drill down performance workflows to troubleshoot bottlenecks	~	~	
	Explore virtualization relationships	Storage and Server virtualization	Storage virtualization	
	Explore replication relationships	~	~	
	Retain performance data	Customizable	1 year	

Table 23. Features in IBM Spectrum Control and IBM Storage Insights for IBM Spectrum Control (continued) **Resource Management Features** IBM Spectrum Control **IBM Storage Insights** (Advanced edition) for IBM Spectrum Control Service Deployment method On-premises Off-premises (in IBM Cloud) Filter Call Home events to quickly isolate trouble spots Hassle-free log collection Simplified ticketing Show active PMRs and ticket history Active directory and LDAP integration for managing users Reporting Inventory, capacity, performance, and storage consumption reports Rollup reporting REST API Alerting **Predictive Alerts** Customizable, multiconditional alerting, including alert policies **Analytics** Performance planning Capacity planning Business impact analysis (applications, departments, and groups) Provisioning with service classes and capacity pools Balance workload across pools Optimize data placement with tiering Optimize capacity with reclamation Transform and convert volumes

Table 23. Features in IBM Spectrum Control and IBM Storage Insights for IBM Spectrum Control (continued)				
Resource Management	Features	IBM Spectrum Control (Advanced edition)	IBM Storage Insights for IBM Spectrum Control	
Pricing		On-premises licensing	No charge for IBM Spectrum Control customers	

Replication products

IBM Spectrum Control no longer supports Tivoli Storage Productivity Center for Replication. IBM Copy Services Manager is the replacement product for replication to use with IBM Spectrum Control.

Information about Copy Services Manager

It is strongly recommended that you back up your replication environment before you proceed with any upgrade.

See "Backing up a replication environment" and verify that your replication backup file location is outside of the IBM Spectrum Control installation directory structure or the replication backup file is deleted during the uninstall of replication.

Copy Services Manager provides automation for 2-site replication, 3-site replication, and advanced copy services from a single point of control across multiple platforms. To learn more about how to download and install Copy Services Manager, see 📑 IBM Copy Services Manager.

Downloading the software code and the license for Copy Services Manager

Before you migrate to Copy Services Manager, obtain the software code and the license.

To download the software code for Copy Services Manager, go to IBM Fix Central at https:// www.ibm.com/support/fixcentral.

To download the license for Copy Services Manager, follow these steps to access the IBM Passport Advantage Online download portal. The license is part of the IBM Spectrum Control eAssembly in IBM Passport Advantage.

Note: You need an IBM ID and password to access Passport Advantage Online for customers.

- 1. Click https://www.ibm.com/support/docview.wss?uid=ibm10874924.
- 2. Scroll to Download Package.
- 3. Click Downloading IBM Spectrum Control Vvrm using Passport Advantage Online (where vrm=the version, release, and modification of the software).
- 4. Scroll to the **Optional parts**, and locate the part number for Copy Services Manager. Make a note of the part number.
- 5. Enter https://www.ibm.com/software/passportadvantage/pacustomers.html to go to Passport Advantage Online for customers.
- 6. Click Customer sign in to enter your IBM ID and password. If you are not an authorized user, follow the steps on the page to request access.
- 7. On the Software and services online page, click Software download & media access.
- 8. Use the **Find by part number** search option, and enter the part number that you copied in Step 4 to download the license.

Migrating to Copy Services Manager

Before you migrate to Copy Services Manager, be sure to back up your configuration of Tivoli Storage Productivity Center for Replication. Then, install Copy Services Manager and uninstall Tivoli Storage Productivity Center for Replication.

To perform the migration, follow these steps:

- Back up your Tivoli Storage Productivity Center for Replication configuration by using the steps in the topic "Backing up a replication environment." The topic is located in the IBM Spectrum Control 5.3.3 Knowledge Center or in the IBM Spectrum Control 5.3.3 Installation Guide. You can find and download the guide from the "Printable documentation" topic in the IBM Spectrum Control 5.3.3 Knowledge Center.
- 2. Download the software code for Copy Services Manager (go to IBM Fix Central at https://www.ibm.com/support/fixcentral and follow the steps).
- 3. Download the license for Copy Services Manager (see the preceding section **Downloading the software code and the license for Copy Services Manager**).
- 4. Install Copy Services Manager. For instructions, see Installing on distributed systems.
 - **Note:** During the installation, you will be asked to choose whether to "Migrate or restore by using existing backup." For details, see Migrating on distributed systems.
- 5. Uninstall Tivoli Storage Productivity Center for Replication by using the steps in the topic "Manually uninstalling Tivoli Storage Productivity Center for Replication." The topic is located in the IBM Spectrum Control 5.2.12 Knowledge Center or in the IBM Spectrum Control 5.2.12 Installation Guide. You can find and download the guide in the "Printable documentation" topic in the IBM Spectrum Control 5.2.12 Knowledge Center.

Product updates and security fixes

IBM Spectrum Control provides regular maintenance updates that can include code fixes, security fixes, new features, and enhancements.

Releases and downloads

IBM Spectrum Control provides regular maintenance updates that can include code fixes, security fixes, new features, and enhancements.

Three types of maintenance packages are available for IBM Spectrum Control:

- **Releases** represent a significant update to the product, which includes, but is not limited to, new features, enhancements to existing features, code fixes, and security fixes.
- **Refresh packs** are quarterly deliverables that update the modification level of a release, and include new features, enhancements to existing features, code fixes, and security fixes.
- Fix packs update the fix level of a release, and primarily include code and security fixes.

All types of maintenance packages are cumulative: a fix pack includes all previous fix packs; a refresh pack includes all the previous fix packs and refresh packs; a release includes all previous fix packs, refresh packs, and releases.

With each new maintenance package that is delivered, the product number changes. The following table provides example of product number changes when you install a package.

Maintenance packages	Product version change examples
Release	5. 3
Refresh pack (modification level)	5.3. 1
Fix pack	5.3.1. 1

A fix pack is installed on top of a specific modification level and release. If you install a fix pack that belongs to a newer level than your current fix pack, then your environment is automatically upgraded to the newer fix pack. For example:

- If you are on fix pack 1 (5.3.1.1) and you apply fix pack 2, then the version is changed to 5.3.1.2
- If you are on fix pack 1 (5.3.1.1) and you apply fix pack 3, then your environment is automatically updated to fix pack 3 and your version is changed to 5.3.1.3.

Tip: When you install a release, the modification level and fix pack number are reset to **0**.

Downloading IBM Spectrum Control

If you have an IBM ID and have purchased IBM Spectrum Control, you can download packages for a new installation or download packages that you can apply as an upgrade or fix to your existing installation.

- To download the package for a new installation of IBM Spectrum Control, go to <u>Passport Advantage</u> Online.
- To download the package for upgrading or applying the latest fixes to your existing installation of IBM Spectrum Control, go to Fix Central.

Tips:

- If you are downloading the package for a new installation, a license is required and is included on Passport Advantage Online.
- If you are downloading the package for an upgrade, your existing license will be transferred.
- Fix Central is the recommended site for retrieving fixes for most IBM products. Fix Central enables you to search, select, and download appropriate fixes for their situation.

Security

In IBM Spectrum Control, the security of your data is critical. At IBM, we're serious about keeping that data safe and constantly strive to address vulnerabilities before they impact your business.

Maintenance and security fixes are regularly provided for the IBM Spectrum Control base product. The base product includes core technology and components such as WebSphere Application Server Liberty profile and the Java Runtime Environment.

When security issues or vulnerabilities are discovered in IBM Spectrum Control or its core components, IBM analyzes the product and distributes a fix or mitigation instructions, as appropriate. These issues are communicated through security bulletins on ibm.com and through discrete communication when necessary.

To find security bulletins, go to the following pages on ibm.com:

- Security bulletins for IBM Spectrum Control
- IBM Support portal
- Supported storage products

Managing security for related software

An IBM Spectrum Control installation also includes other products, such as Db2. You can monitor fixes for these related products in the same way that you do for IBM Spectrum Control on ibm.com:

- IBM support portal for Db2: https://www.ibm.com/support/entry/myportal/product/information_management/db2_for_linux,_unix_and_windows
- IBM support portal for IBM WebSphere Application Server: https://www.ibm.com/support/entry/ myportal/product/websphere_websphere_application_server
- IBM Product Security Response Team blog at https://www.ibm.com/blogs/psirt/

Each product team is responsible for publishing their own security bulletins. Review any published security bulletins to identify affected versions and follow the remediation or fix information. If these products are affected, they might provide a maintenance fix. It is acceptable to apply a fix pack to these products to resolve security vulnerabilities or other problems.

Important: Upgrading to a new version or release of these products is not typically supported by IBM Spectrum Control and might result in compatibility issues. Even with fix pack updates, read the security bulletins carefully for any compatibility concerns. For example, if a common protocol is disabled, both the related product and IBM Spectrum Control must also accommodate a new protocol.

Subscribing to IBM announcements

Subscribe to My Notifications to be notified automatically of IBM announcements such as security bulletins and flashes for IBM Spectrum Control.

With My Notifications, you can specify that you want to receive daily or weekly email announcements. You can specify what type of information you want to receive (such as publications, hints and tips, product flashes (also known as alerts), downloads, and drivers). You can also customize and categorize the products about which you want to be informed and the delivery methods that best suit your needs.

- 1. Go to http://www.ibm.com/software/support/einfo.html.
- 2. Click Subscribe now!.
- 3. Log in with your IBMid.
- 4. In Product lookup, type IBM Spectrum Control.
- 5. Click **Subscribe** for the edition of IBM Spectrum Control that you use.
- 6. Select the types of documents for which you want to receive notifications and click **Submit**.

End of support

End of Support (EOS) marks the official withdrawal of technical support for specific versions and releases of IBM Spectrum Control and other IBM software products.

What does End of Support mean to you

Important:

The following releases have reached end of support:

- Tivoli® Storage Productivity Center 5.2.0 5.2.8, 5.2.8 5.2.17
- IBM Spectrum Control 5.2.8 5.2.17

When a specific version or release of IBM Spectrum Control reaches end of support (sometimes called end of life), you might find that access to entitled severity 1 technical support is no longer available. The documentation for each product version is still available in PDF format when you access its IBM Knowledge Center.

What can you do to monitor End of Support dates

Review the End of Support Announcement letters regularly, see https://www.ibm.com/software/passportadvantage/endofsupportannouncementletters.html

End of support announcements are generally made twice a year in April and September and are subject to change. You can confirm dates by visiting the IBM Software Support Lifecycle policy site, see https://www.ibm.com/support/home/pages/lifecycle/index.html

For a complete list of supported versions of IBM Spectrum Control, see https://www.ibm.com/support/pages/overview-ibm-spectrum-control-releases

End of support options

The best option for you when a version or release of IBM Spectrum Control is going into end of support is to upgrade to a newer version or release. You not only maintain access to new fixes and features, but you can open support cases that ensure if you do have an issue that requires immediate attention that you can get it.

You can purchase an IBM Software Support Service Extension (also known as Extended Support). Service Extensions give you access to product fixes and the ability to engage with IBM Support for use and defect support after your End of Support date is reached. Note, these extensions are not available for all products. Ask your IBM Support representative if IBM Spectrum Control is eligible for the extension.

Collaborating with the team

Collaborate with the IBM Spectrum Control team to help improve the product.

Got a great idea for making IBM Spectrum Control even better? Do you want to vote for an enhancement that was requested by another user? The IBM RFE Community is a place where you can collaborate with the development team for IBM Spectrum Control and other product users through your ability to search. view, comment on, submit and track product requests.

- 1. Go to the Servers and Systems Software RFE community at https://www.ibm.com/developerworks/rfe/ execute?use_case=changeRequestLanding&PROD_ID=420.
- 2. To submit an enhancement request, click the **Submit** tab.
- 3. Log in with your IBMid.
- 4. View existing requests to ensure that your idea hasn't already been submitted. A link to the search page is included in the form.
- 5. Complete the form and click **Submit**.

Reporting a problem

Report problems that you encounter in IBM Spectrum Control.

IBM is committed to customer satisfaction and is ready to assist if you encounter any problems with IBM Spectrum Control. By opening a case, you can ensure that IBM Support will be informed of your problem. When you open a case, have the following information ready:

- The version, release, modification, and service level number of your IBM Spectrum Control installation.
- 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.

To open a case for IBM Spectrum Control, complete the following steps:

- 1. Go to the IBM Support portal.
- 2. Click Open a case.
- 3. Log in with your IBMid.
- 4. Describe the problem.
- 5. Submit the case.

Videos: The IBM Support portal is supported by IBM Watson and provides you with enhanced transparency into ticket resolution workflow and improved self-service options. Watch a few short videos to learn more:

- Introducing A New Customer Portal
- Open And Manage Cases
- IBM Support Community: Search
- IBM Support Community: Forums

Key concepts

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

Data collection

To help you implement a storage management strategy, it is critical that you determine the information that you want to gather about the resources within your environment. You can schedule different data collection jobs depending on the types of information that you want to gather and the resources that you want to monitor.

Use probes and performance monitors to collect detailed information about resources. You must schedule and run these data collection jobs before you can complete other tasks within IBM Spectrum Control, such as viewing reports, setting alerts, and managing storage systems.

You schedule probe and performance monitor jobs for a resource when you add the resource for monitoring. After you add the resource, you can modify the probe and performance monitor job schedules on the list and details pages for the resource. For example, to modify the probe job schedule for a storage system, go to the **Storage Systems** page or the **Storage System details** page in the GUI.

Schedule data collection jobs to gather the following types of information:

Asset and status information

Use <u>probes</u> to collect asset, storage statistics, and status information about resources. You can run probes on all the resources that are monitored by IBM Spectrum Control.

Performance monitoring

Use <u>performance monitors</u> to collect metrics that measure the performance of switches and storage systems.

Performance monitoring and troubleshooting

IBM Spectrum Control can collect information about the performance of storage systems and switches. This information includes key performance metrics and alerts of threshold violations that can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage.

To monitor the performance of resources and check for threshold violations, complete the following tasks:

- Add resources for monitoring and schedule data collection
- · Define alerts for performance thresholds
- · View and troubleshoot performance issues

Collect performance data

Before you can troubleshoot and view reports about performance, you must collect data about monitored resources. Performance monitors are data collection jobs that gather performance information about resources. This information includes metrics that measure the performance of the components within a resource. Metrics measure the performance characteristics of volumes, ports, and disks on storage systems and switches. IBM Spectrum Control provides many different metrics for measuring performance. For example, some key metrics for storage systems are I/O rate in I/O operations per second, data rate in MiB per second, and response time in milliseconds.

You can use metrics in IBM Spectrum Control to track growth or change in I/O rates, data rates, and response times. In many environments, I/O and data rates grow over time, and response times increase as those rates increase. This relationship can help with "capacity planning" for your storage. As rates and response times increase, you can use these trends to project when more storage performance and capacity is required.

Define alerts for performance thresholds

Alerts can notify you when the performance of a monitored resource falls outside of a specified range and might represent a potential problem. When you define an alert for an internal component of a resource, select a specific metric that you want to measure and its boundary values. When the performance of a resource falls outside the boundary values, an alert is triggered.

For example, you can define an alert that is triggered when the overall back-end response time for a managed disk on a SAN Volume Controller exceeds a certain value. The overall back-end response time is a metric that measures the average number of milliseconds that it takes to service each I/O operation on a managed disk.

View and troubleshoot performance issues

After data collection and performance thresholds are configured, you can use the web-based GUI to complete the following tasks:

- Measure, compare, and troubleshoot the performance of switches, storage systems, and their internal resources.
- Review the threshold violations and alerts that were triggered when the performance of a resource fell outside of a specific range.
- View performance information in a chart or table format to help you quickly identify where and when performance issues are occurring. The chart is a visual representation of how the performance of resources trend over time.
- Customize views of performance so that you can analyze specific resources and metrics during time ranges that you specify.
- Drill down into resources to view detailed information about the performance of internal and related resources. For example, if a SAN Volume Controller storage system is shown in the chart, you can quickly view and compare the performance of its internal and related resources, such as disks, volumes, ports, managed disks, and back-end storage.
- Implement server-centric monitoring of SAN resources without requiring a Storage Resource agent. When you add an agentless server for monitoring, IBM Spectrum Control automatically correlates that server with the ports on known host connections. If matches are found between the server and host connections on monitored storage systems, you can view the performance of the internal resources that are directly associated with the SAN storage that is assigned to the server. For example, if a SAN Volume Controller maps two volumes to the server, you can view the performance of those volumes and the related managed disks.
- Export performance information to a CSV file. A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application.
- In the optional Cognos Analytics reporting tool, you can also view and create performance reports about multiple resources.

Related concepts

"Monitoring the performance of resources" on page 260

IBM Spectrum Control can collect information about the performance of storage systems and switches. This information includes key performance metrics that can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage.

Related tasks

"Adding resources" on page 115

You must add resources to IBM Spectrum Control before you can collect data, generate reports, and manage storage that is related to those resources.

Related reference

"Performance metrics" on page 277

IBM Spectrum Control can collect information about the performance of storage systems and switches. This information includes metrics that measure the performance of volumes, ports, and disks. You can

view performance information or specify alerts for when the performance of a resource falls outside a specified range.

Applications, departments, and general groups

Use applications and departments to organize the resources in your storage environment into a hierarchy that matches the structure of your business organization. The information can be used to view the performance of resources that belong to an application and the capacity growth of a department. Use general groups to group your resources, such as the storage systems with lease agreements that end in the current year, so that you can view information about the resources at one location in the GUI.

IBM Spectrum Control provides storage grouping to define resources that utilize application and department business modeling. The information can be used to view the performance of resources that belong to an application and the capacity growth of a department. The pages for the monitoring and management of applications and departments are aligned with user scenarios, for example, capacity trending, resource health monitoring and performance troubleshooting

To manage and administer applications and departments, complete the following tasks:

- Create applications and departments to model data for storage resources, allowing for enhanced capacity trending and performance troubleshooting.
- View the status of resources that make up the applications and departments.
- Create resource filters to automatically add resources to applications.
- Add resources directly to applications.
- Add applications to existing applications and departments.
- Add departments to existing departments.
- View subcomponent (member application) information to support hierarchical levels of storage grouping.

Applications

An application is a program or a project that consumes storage resources within an organization and interfaces with other enterprise groups that are important to the running of a business. Use IBM Spectrum Control to model the storage usage that is consumed in your environment by assigning the usage to applications to see the overall health status.

An application can be part of a department and have its own subcomponents (member applications) that are used to create a five level deep hierarchy. Applications that are grouped together can range from large line-of-business systems to specialized software, in a department, that runs on either client computers or servers. For example, an application might be an automated billing system within the Finance department, VMware running in the Information Technology department or an email marketing system that is part of the Marketing department.

The Applications page shows capacity information about the applications that are monitored by organization and interfaces with other enterprise groups that are important to the running of a business. Use IBM Spectrum Control to view and manage the resources that are defined to the applications. If the application is associated with a department, information about that department is also shown.

Departments

A department is a division within a business. Use IBM Spectrum Control to model the storage capacity that is consumed in your department for your business environment, in accordance with other department members.

A department can be hierarchical in its organizational layout. For example, a department might use 15 applications and be part of another five departments. A department might share storage resources with another department, subdepartment or an application even if they don't belong in the same hierarchy. For example, in a collaboration scenario, a single IBM SAN Volume Controller might be shared by multiple departments.

The Departments page shows storage capacity information about the top-level departments, the subdepartments and any applications that belong to the department that are monitored by IBM Spectrum Control.

General groups

Create general groups to quickly view information about storage resources that have common characteristics. For example, you might group the subset of ports on a SAN Volume Controller that are used for inter-node communication or the storage systems that are used by a critical business application.

You can organize your storage resources into a general group hierarchy. Organizing resources into general groups and their subgroups can be helpful when you want to quickly view information about a group of resources, but you also want to view information about subgroups of resources within the group. For example, you can group the resources that are used by your production application so that you can monitor all the application resources and separately monitor the specific subgroups of storage systems and ports.

Storage optimization

To optimize the resources in your storage environment, you can create tasks to balance pools, re-tier volumes, or transform volumes that are in storage virtualizer pools.

Discontinued support: Storage optimization is no longer supported in IBM Spectrum Control. While the feature might still work in this release, it's recommended that you use another tool for your optimization needs, when possible. For a complete list of discontinued features, see "Discontinued features in IBM Spectrum Control" on page xxx.

Balancing the workload of volumes across storage pools

You can balance the workload of volumes across pools on the same tier. The pools are analyzed and recommendations are generated to move volumes from pools with high-activity values to pools with low-activity values.

Re-tiering volumes

You can re-tier volumes to balance pools or to tier volumes based on the criteria that you set in tiering policies. For example, you can tier volumes that are based on the volume workload or on file usage, or both. Depending on the conditions that are set in the tiering policy, recommendations are generated. For example, you can reduce storage costs by moving volumes with low workloads to lower or less expensive tiers. You can also improve performance and use storage more efficiently by moving volumes with heavy workloads to the tiers that best meet their workload requirements.

Transforming volumes

You can complete the following tasks for one or more volumes in storage virtualizer pools:

- Move volumes from one storage virtualizer pool to another pool on the same storage virtualizer.
- Move volumes in a storage virtualizer pool to a pool that is enabled for Easy Tier.
- Convert fully allocated volumes to thin-provisioned volumes and convert thin-provisioned volumes to fully allocated volumes.
- Convert fully allocated volumes to compressed volumes and convert compressed volumes to fully allocated volumes.

The pools are analyzed and recommendations are generated. For example, analysis occurs to ensure that there is sufficient space in the pool to convert volumes or add volumes.

Related tasks

"Optimizing storage pools" on page 444

Analyze the activity of pools and resolve performance hot spots by redistributing volumes across each storage tier.

"Optimizing storage tiering" on page 438

To optimize the placement of volumes on storage tiers, analyze the tiering of volumes in your storage environment.

"Transforming and migrating volumes" on page 448

You can transform fully allocated volumes to compressed or thin-provisioned volumes, or transform compressed or thin-provisioned volumes to fully allocated volumes. You can move volumes to other pools or to pools that are enabled for Easy Tier.

"Managing tasks for tiering storage, balancing pools, and transforming storage" on page 461
To optimize the resources in your storage environment, use tiering storage, balancing pools, and transforming-storage tasks. Use the **Tasks** page to manage all the tasks that are used to optimize storage. Use resource list and resource detail pages to manage the optimization tasks for resource types and for specific resources.

Related reference

"Criteria for identifying the pools that require balancing" on page 446

To identify the pools that require balancing, performance data is collected to determine the activity of pools on the same tier. The performance capability of pools is preserved by setting an activity limit for each tier of storage.

Cloud configuration

To take advantage of the simplified provisioning and optimization capabilities in IBM Spectrum Control, some configuration is required.

Discontinued support: Cloud configuration, provisioning, and optimization are no longer supported in IBM Spectrum Control. While these features might still work in this release, it's recommended that you use another tool for your provisioning and optimization needs, when possible. For a complete list of discontinued features, see "Discontinued features in IBM Spectrum Control" on page xxx.

To provision storage, you are required to specify only the storage capacity and storage quality that is required. After volumes are created, IBM Spectrum Control can analyze and optimize volume performance. To take advantage of these capabilities, create service classes and, optionally, capacity pools.

A tutorial of cloud configuration concepts is available in the GUI. In the menu bar, go to **Advanced Analytics** > **Cloud Configuration**, and click **Learn the Concepts**.

Service classes

A service class is a set of properties that describe capabilities and characteristics of storage resources. A service class typically describes a particular quality of service, and is used during provisioning to describe storage requirements. For example, a block-storage service class specifies properties such as a required RAID level, and whether storage resources must be able to encrypt or thin provision volumes.

Service classes simplify provisioning requests by representing a level or type of storage quality. When you are requesting storage, only the required capacity and service class must be specified. Before you can provision storage, you must create service classes that describe the capabilities and characteristics of the storage you want to be able to provision. Service classes can later be modified or deleted as the needs of your installation change.

Block-storage service classes

A block-storage service class describes attributes and capabilities of block storage resources. When you provision volumes, you specify the requirements by using a block-storage service class. Based on the requirements of the service class, IBM Spectrum Control identifies a storage pool for the volume.

File-storage service classes

A file-storage service class describes attributes of file storage resources. The file storage resources include Network Attached Storage (NAS) filers, Network Shared Disks (NSDs), and file systems. When you provision shares, you specify the requirements by using a file-storage service class. Based on the requirements of the service class, IBM Spectrum Control identifies a file system or NSD for the share. Space on an NSD is allocated if a new file system is required or an existing file system requires more space.

Service class properties

A service class represents storage resources with common traits. These common traits are the service class properties, and collectively describe a particular quality of service. Before you create service classes, consider the levels of storage quality that you want to have available for provisioning requests. For block storage, for example, you can organize service classes by storage tier or RAID level. Some service class properties are used during provisioning to identify the best location for storage placement. Other service class properties determine how the storage is configured, and reflect the performance, reliability, or security considerations for the level of service.

In addition to the standard properties of a service class, you can create your own properties by using custom tags. Custom tags can represent any common trait that storage systems must have. For more information about creating your own properties by using custom tags, see "Custom tags" on page 70.

Block-storage service classes can specify the following properties:

- Storage tier or range of storage tiers
- RAID level
- Whether volume virtualization is required or not allowed. This property is available only when at least one IBM storage virtualizer is managed by IBM Spectrum Control. The IBM storage virtualizers include SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified.

If volume virtualization is required, the VDisk mirroring property specifies whether volumes are mirrored to a second storage pool, which uses different backend storage, on the same IBM storage virtualizer.

• Whether volumes must be thin provisioned or must not be thin provisioned.

When a thin-provisioned volume is created in a DS8000, an extent space efficient (ESE) volume is created.

If thin provisioning is required, you can specify thin-provisioning configuration properties for the IBM storage virtualizers and IBM XIV Storage System. The following properties might be available:

- If the volume is created on an XIV, the locking behavior configuration property affects the volume that is created.
- If the volume is created on an IBM storage virtualizer, the following configuration properties affect the volume that is created:
 - Used capacity
 - Auto-expand
 - Warning level
 - Grain size
- Whether volumes are compressed. This property is applied only when an IBM storage virtualizer has at least one I/O group that contains a compressed volume. If there are no existing compressed volumes in an I/O group, IBM Spectrum Control will not create compressed volumes.

This property is available only when the thin provisioning property is enabled.

• A multipathing policy. If a server has a Storage Resource agent (SRA) and uses an IBM System Storage Multipath Subsystem Device Driver (SDD), the server multipathing policy is used to configure the driver.

If a multipathing policy is set, and fabrics are managed by IBM Spectrum Control, the following properties affect how the storage system is connected to the host:

- Whether the storage system and the client host are connected through fully redundant paths
- Number of paths between the storage system and the client host
- If at least one DS8000 Encryption Group is discovered by IBM Spectrum Control, the following properties are available and affect the storage placement and configuration of volumes:
 - Whether volumes must be encrypted
 - If volumes must be encrypted, the DS8000 encryption group to use

File-storage service classes have the following properties:

- Whether dedicated storage is required, or whether shared storage is allowed. In other words, whether the file system from which the NAS share is provisioned can contain other NAS shares.
- Fileset type. The fileset type is either an independent fileset with its own allocated nodes, or a dependent fileset allocated on the file system.
- · Access path host name replacement.

Predefined service classes

To illustrate service classes, a set of predefined service classes are provided. From the **Service Classes** page in the web-based GUI, you can view any of these service classes for more information. Select a service class and then select **Actions** > **View / Modify**. To display help information about a service class property, complete the following steps:

1. Move the mouse pointer over the field or control.

A question mark icon is displayed next to the field.

2. Move the mouse pointer over the question mark icon to display an explanation of the property.

The following sample block-storage service classes are provided for illustration.

Gold

The Gold service class represents the highest-performing storage resources for mission-critical applications.

Silver

The Silver service class represents high-performing storage resources for applications in production.

Bronze

The Bronze service class represents standard storage resources for non-mission-critical applications.

The following sample file-storage service classes are provided for illustration.

NormalIsolation

The NormalIsolation service class describes normal isolation file storage by specifying that shared storage is allowed. The file system from which the NAS share is provisioned can contain other NAS shares.

EnhancedIsolation

The EnhancedIsolation service class describes enhanced isolation file storage by specifying that dedicated storage is required. The file system from which the NAS share is provisioned cannot contain other NAS shares.

Custom tags

During provisioning, IBM Spectrum Control determines candidates for storage placement by comparing the requirements of a service class with the known capabilities, configuration, and performance of the available storage systems. You can think of the storage placement determination as a filtering process. IBM Spectrum Control filters the set of available storage resources against properties of the service class until only the storage resources that can satisfy all of the requirements remain.

In addition to the standard properties of a service class, you can create custom requirements for the service class by specifying up to three custom tags. To provide the service class, storage resources must have all the same tags that are specified in the service class.

When you specify tags on a block-storage service class, only pools that have all the same tags are candidates for provisioning. If a pool is not tagged, any tags on the containing storage system also apply to the pool.

When you specify tags on a file-storage service class, only file systems and NSDs that have all the same tags are candidates for provisioning. If a file system or NSD is not tagged, any tags on the containing storage system also apply to the internal resource.

During provisioning, IBM Spectrum Control filters candidates for provisioning against the custom tags.

Before you create service classes, consider whether there are any special requirements for provisioning that are not addressed by the standard properties of the service class. If so, tag the appropriate storage resources to satisfy custom requirements of a service class. When you are defining the service class, specify the custom tags that are required.

Adding users to a service class

By default, only administrators can provision storage. However, an administrator can add users to a service class to grant them permission to provision by using the service class. In order for an administrator to grant permission to provision to a user, the user must be assigned to the Monitor or External Application role.

When you are granting users permission to provision by using a particular service class, you can also specify whether execution of provisioning plans created by these users requires administrator approval. The requirement for administrator approval applies to all non-administrative users who are added to the service class, and also applies to external applications. Approval cannot be required for some non-administrative users and not for others.

Before you create service classes, consider whether you want to allow users to request their own storage. Consider whether you want user provisioning requests to require administrator approval.

Associating capacity pools with a service class

Before you create service classes, consider whether certain storage requests must always be satisfied from a particular set of resources. If so, you can associate one or more capacity pools with a service class. For example, suppose that a set of storage resources are allocated to a particular department of your business. All storage requests for that department must be satisfied by those storage resources. In this case, you can add all the storage resources to a capacity pool and associate that capacity pool with a service class. You can also add the members of the department to the service class. The department members can then request their own storage as needed. However, the requests can be satisfied only by the storage resources that are allocated to the department.

When you are configuring capacity pools and service classes, keep in mind that a storage resource can be a member of only one capacity pool. Adding storage resources to a capacity pool is not a viable approach if the resources are already members of another capacity pool. Consider specifying custom tags to achieve the same result. In the preceding example, suppose that some or all of the storage resources that are allocated to the department are already members of capacity pools. In this case, you can instead tag the storage systems and the service class with matching tags to achieve the same result. With either approach, only the storage systems that are allocated to the department are candidates for provisioning when storage is requested by using the service class.

Capacity pools

Capacity pools are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. Configure capacity pools to track the used and available capacity for block and file storage on any set of storage resources. Provisioning requests can also be restricted to resources in a capacity pool.

The following types of storage resources can be grouped into capacity pools:

- · Storage systems
- · Storage pools
- File systems of file storage systems

You can organize storage resources into capacity pools in any way that serves your business needs. For example, you can separate the storage resources that are allocated for separate divisions of your business into separate capacity pools. You can then track the storage use for each division separately, and restrict provisioning requests to the appropriate set of storage resources.

Tracking used and available capacity in a capacity pool

You can track the used and available capacity for any set of storage resources by adding the storage resources to a capacity pool. From the Capacity Pools page of the GUI, you can view the total capacity of the resources, and monitor the following capacity measurements:

- · Used file space
- · Available file space
- Used block space
- · Available block space

Restricting a provisioning request to a capacity pool

You can use capacity pools to define a set of storage resources from which provisioning requests must be satisfied. When you are provisioning storage, you can specify a capacity pool. If you do, the provisioning request is restricted to resources in the capacity pool. Only those resources are candidates for provisioning.

You can also associate a service class with capacity pools. If a service class is associated with capacity pools, provisioning requests for the service class must specify, and be constrained to, one of the associated capacity pools. Because you can also grant non-administrative users permission to provision storage by using the service class, associating capacity pools with the service class restricts the users to a specific set of resources.

Provisioning storage overview

IBM Spectrum Control guides you through the steps for provisioning storage volumes or network-attached storage (NAS) shares. You can provision storage volumes or NAS file shares to one or more servers, one or more hypervisors, or one cluster.

Discontinued support: Provisioning is no longer supported in IBM Spectrum Control. While the feature might still work in this release, it's recommended that you use another tool for your provisioning needs, when possible. For a complete list of discontinued features, see "Discontinued features in IBM Spectrum Control" on page xxx.

Storage requirements

When you set up provisioning, you must specify your storage requirements by defining a service class. Certain properties of the service class describe capabilities that storage resources must have so they can provide the service class. For example, when you provision volumes, a block-storage service class specifies properties such as a storage tier, a RAID level, and whether the pool must be able to encrypt or thin provision volumes.

A service class typically represents a quality or service. For example, IBM Spectrum Control provides predefined block-storage service classes that are named Gold, Silver, and Bronze. Gold represents the highest-performing storage resources, and Silver and Bronze represent lesser levels of storage quality.

The different levels of storage quality that are provided by each service class are defined by the service class properties. For example, each service class specifies a different required storage tier that depends on the level of service that is required. Because the Gold service class is for mission critical applications, it sets off thin provisioning.

When IBM Spectrum Control guides you through the steps for provisioning, you concern yourself only with which service class is needed for the new volumes or shares, and how much capacity you require. However, if you are unfamiliar with the service classes, you can open a separate window to view service class information. If you have Administrator privileges, you can modify or create service classes.

Storage constraints

IBM Spectrum Control identifies the storage resources that can provide the capacity and the service class from a set of storage resources. This set might be all the storage resources that are known to IBM Spectrum Control, or it might be constrained to a subset of those resources. In particular, your site can create capacity pools. *Capacity pools* are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. For example, a capacity pool might contain the storage resources that are allocated to a particular department or division of your business.

A service class might restrict storage placement to one or more capacity pools. If the service class you are provisioning from restricts placement, you must select one of the allowed capacity pools. If the service class does not restrict storage placement, you can specify that candidates for provisioning can be selected from all storage resources that are known to IBM Spectrum Control, all storage resources in any capacity pool, or storage resources in a particular capacity pool. If you are unfamiliar with the capacity pools that you can select when you set up provisioning, you can open a separate window to view capacity pool information. If you have Administrator privileges, you can modify or create capacity pools.

If the fabrics are managed by IBM Spectrum Control, the set of storage resources might be further constrained by your zoning policy. If the zoning policy is not configured to automatic zoning, only storage systems with connectivity to the server or hypervisor are candidates for storage placement.

Provisioning volumes

Restriction: To provision volumes, you must have the IBM Spectrum Control Advanced Edition license.

To provision volumes, you must select one or more servers or hypervisors. You can request one or more volumes. For each volume, you specify a volume name, the capacity that is required, a service class, and, optionally, a capacity pool. For each volume, IBM Spectrum Control identifies the storage pools that can provide the capacity and the service class. From the set of pools that can provide the capacity and service class, IBM Spectrum Control identifies the best location for the storage. The best location for the storage is based on the available capacity in the pool and performance data. Preference is first given to storage pools and systems that already contain volumes for the selected server or hypervisor. Preference is then given to systems that have available performance data.

On storage systems, volumes are allocated in increments of a set unit size that depends on the storage system type. Depending on the storage system that was identified as the best location for the storage, the volumes that are allocated might be larger than the capacity you requested for the following reasons:

- On SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified, volumes are allocated in increments of an extent size. The extent size is set when the MDisk group is created. The capacity that you requested is rounded up, if necessary, to the full extent size.
- On a DS8000, volumes are allocated in increments of a 1 GiB fixed extent size. The capacity that you requested is rounded up, if necessary, to the full extent size.
- On an XIV, volumes are allocated in fixed increments of 16 GiB. The capacity that you requested is rounded up, if necessary, to complete the 16 GiB increment.

Provisioning shares

To provision a share, you must first select one or more servers or hypervisors. You specify the capacity that is required for the share, a service class, and information about how to export the share. IBM Spectrum Control identifies the file systems and Network Shared Disks (NSDs) that can provide the capacity and the service class. From the set of file systems and NSDs that can provide the capacity and service class, IBM Spectrum Control identifies the best location for the storage. The best location for the storage is based on the available capacity on the file system or NSD.

The provisioning task

When you complete the steps for setting up provisioning, IBM Spectrum Control creates a provisioning task for the share or volume. If you request multiple volumes, a separate task is created for each unique service class and capacity pool combination you specify in your volumes request. If a requested volume is not constrained to a capacity pool, the set of all available storage resources is considered the capacity pool. If you have Administrator privileges, you can save, run, or schedule provisioning tasks.

In the service class, an administrator can grant users permission to provision by using the service class. In the service class, an administrator can also specify whether scheduling or running provisioning tasks that are created by using the service class requires administrator approval. If you are a user with permission to provision by using the service class, and administrator approval is not required, you can save, run, or schedule the task. You can save the task, if administrator approval is required. Also, if you save the task, an administrator can later schedule or run it.

Although IBM Spectrum Control identifies the best location for storage when it creates the provisioning task, the implementation of the task might fail. For example, the implementation of the task might fail because of changes in the environment that occur after the task is created and before it is run. The implementation of the task might also fail because of storage system restrictions that are not considered when the task is created.

Viewing task details and logs

When a provisioning task is running, you can view its status in a details page. If you run the provisioning task immediately after you complete the steps for setting up provisioning, the details page is already displayed. You can also display details of a provisioning task from the **Tasks** page. From the details page, you can open logs for the provisioning task. The logs show the steps that are taken by the task during processing and include detailed information about any warnings or errors.

Tracking

When you set up provisioning, you can specify a ticket identifier for tracking purposes. The ticker identifier is associated with the provisioning task, and with any volume or share that is created by the provisioning task.

The ticket identifier can be viewed by showing the Ticket column in the **Volumes** page or the **Shares** page. The ticket identifier can also be viewed in the properties notebook for a volume or share.

Tip: If you request multiple volumes, a separate provisioning task is created for each unique service class and capacity pool combination you specify in your volumes request. By specifying a ticket identifier when you request the volumes, you can easily track the set of provisioning tasks and volumes that are created by the single request.

Storage configuration

Whether you are provisioning volumes or shares, IBM Spectrum Control configures the storage and resources as specified by certain properties of the service class, and, in the case of volumes, according to your zoning policy.

When you provision volumes, IBM Spectrum Control can configure volumes and resources according to properties of the service class:

- If the service class requires thin provisioning, the volume is created in a pool that can thin provision volumes. During provisioning, other properties of the service class are used to configure the thin provisioned volume for the particular type of storage system that contains the pool. Because you configure thin provisioned volumes differently for the different storage system types, the service class has a unique set of thin provisioning properties for each type.
- If the service class specifies a multipathing policy for servers, the IBM System Storage Multipath Subsystem Device Driver (SDD) on the server is configured to use that policy. For the SDD to be configured, a Storage Resource agent (SRA) must be running on the server.

When you are provisioning volumes, if automatic zoning is enabled, new zones might be created to connect a server to the storage system. Existing zones are used if the server already has connectivity to the storage system. Otherwise, one or more new zones are created between a host initiator port and a controller, node, or module port.

When you provision shares, IBM Spectrum Control configures at least one file access protocol for the file share.

Storage reclamation

You can use volume reclamation recommendations to reclaim unused storage capacity in your environment and use your storage more efficiently.

The block storage systems in your data center that you add for monitoring are regularly analyzed. You can view a list of the volumes that are not mapped to servers or that have no recorded I/O activity. By reclaiming the volumes that are not being used, you can recycle existing storage space instead of purchasing new storage media.

Supported storage systems: Go to see http://www.ibm.com/support/docview.wss? uid=swg21386446, and then click the release number in the **Storage** column.

Related tasks

"Reclaiming storage" on page 446

Use the recommendations to reclaim capacity before you plan new capacity purchases.

Alerts and alert policies

Specify conditions that trigger alerts and the actions to take when those alerts are triggered, such as notify an email address. Use alert policies to define those alert conditions and notification settings for a group of resources.

Alerting functions examine the attributes, capacity, and performance of resources. If the conditions that are defined for alerts are met, the actions that are specified for the alert are taken. Typically, the actions include sending a notification. For example, if the status of a SAN Volume Controller storage system changes to Error, an alert is displayed in the **Alerts** page in the GUI, and an email might be sent to a storage administrator.

You can manage alerts in your storage environment in the following ways:

• Use alert policies to manage the alert definitions and notification settings that apply to different sets of resources. For example, you can use one alert policy for the storage systems in your test environment, and another for the storage systems in your production environment.

Alert policies manage one type of resource only. For example, if you have SAN Volume Controller and FlashSystem 900 storage systems in your storage environment, you cannot have both types of resource in one alert policy.

A resource can be managed by only one alert policy. When you add a resource to be monitored by IBM Spectrum Control, it is added to a default alert policy automatically.

If a resource is managed by a policy, the resource cannot have alert definitions and notification settings that are independent of the policy. The alert definitions and notification settings that apply to the resource come from the policy.

Default policies with alerts already configured are available. You can create copies of the default policies and assign resources to the new policies. Your alerts are configured with the default settings.

- Define alert conditions and notification settings for individual resources. It is not a requirement for resources to be managed by an alert policy. A resource can have its own alert definitions and notification settings, independent of an alert policy.
- Define alerts and notification settings for applications and general groups. Use applications or general groups to manage alerts for groups of resource components such as volumes or pools. For example, you might want to define alerts on the response time for volumes in an application, depending on the response time requirements of the application. In this case, it is not useful to configure volume

response time thresholds for the entire storage system because the storage system might serve many different applications with different needs.

Triggering conditions for alerts

The conditions that trigger alert notifications depend on the type of resource that you are monitoring. In general, the following types of conditions can trigger alerts:

- An attribute or configuration of a resource changed
- The capacity of a resource fell outside a specified range
- The performance of a resource fell outside a specified range
- The storage infrastructure was changed, such as a new or removed resource
- Data is not being collected for a resource

For example, you can use performance thresholds to be notified when the total I/O rate for storage systems falls outside a specified range. This information can help you identify areas in your storage infrastructure that are over used or under used. IBM Spectrum Control provides many metrics for measuring performance and determining violations of the thresholds that you specify.

Alert notifications and triggered actions

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 Netcool/OMNIbus events, login notifications, entries in Windows event log or UNIX syslog, and emails. Additionally, if a Storage Resource agent is deployed on a monitored server, you can run a script or start an IBM Tivoli Storage Manager job in response to the alert.

Acknowledging alerts: Some alerts are triggered by conditions that commonly occur and can be ignored. In such cases, you acknowledge these alerts to indicate that they were reviewed and do not require immediate resolution. By acknowledging alerts, you can more quickly identify other alerts that must be reviewed and resolved.

Event processing

Alerts are generated when particular conditions are detected during data collection and event processing. For some storage systems such as IBM Spectrum Accelerate and the XIV, events are polled every minute from the resource.

For other resources, events are subscription-based, where the resource itself or a data source such as an SMI-S provider (also called CIM agent or CIMOM) sends the events to IBM Spectrum Control when conditions change on the resource. Examples of storage systems that use subscription-based event processing include SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, FlashSystem V840, and FlashSystem V9000. For these storage systems, a probe is automatically run when many events are received from the storage system in a short time period. To avoid performance bottlenecks, probes are run only every 20 minutes.

Prerequisites for using alerts

The following conditions must be met to successfully use alerts:

- Data collection schedules are configured and scheduled to run regularly. For example, to detect violations of performance thresholds, you must run performance monitors to collect performance data about resources. Running performance monitors regularly also helps to establish a history of performance for trending analysis.
- 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, IBM Tivoli Netcool/OMNIbus events, or email, you must configure those alert destinations before you use the alert.
- 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 IBM Spectrum Control to listen to SNMP traps. The default port number is 162, and the default community is public.

Related concepts

"Alerting" on page 151

Specify conditions that trigger alerts and the actions to take when those alerts are triggered, such as notify an email address. Use alert policies to define those alert conditions and notification settings for a group of resources.

Reporting

Use reporting functions to view overview and detailed information about your storage.

You can view detailed capacity information for storage systems, servers, hypervisors and their internal resources, such as volumes, pools, disks, and virtual machines. For example, you can view the total amount of storage capacity that is committed to a volume, the capacity that is used by a volume, and the capacity that is allocated to a volume and is not yet used.

You can configure rollup servers to collect information from the resources they manage and roll that data up to master servers. This gives you a network-wide perspective of storage usage in your environment when you have multiple IBM Spectrum Control servers deployed.

You can use the IBM Cognos Analytics reporting tool to view predefined reports and create custom reports about the resources managed by IBM Spectrum Control. Use the IBM Cognos Analytics reporting tool to view predefined reports about the capacity and performance of your resources. Charts are automatically generated for most of the predefined reports. Depending on the type of resource, the charts show statistics for space usage, workload activity, bandwidth percentage, and other statistics. You can schedule reports and specify to create the report output in HTML, PDF, and other formats. You can also configure reports to save the report output to your local file system, and to send reports as email attachments.

You must collect information about your environment before you can use reports to view details about the storage resources in it. You can use IBM Spectrum Control monitoring jobs, such as probes and performance monitors, to gather comprehensive information and statistics about your storage resources.

Rollup servers

Rollup reporting combines capacity and status information from multiple instances of IBM Spectrum Control for monitoring and reporting. In rollup mode, subordinate servers collect information from managed resources and roll that data up to master servers.

Rollup capability requires two types of servers: a master rollup server running IBM Spectrum Control, which collects capacity and status information from one or more subordinate rollup servers. Subordinate servers communicate with the master server during server probes.

The master server collects capacity and status information about storage resources that are managed by the subordinate servers. This master/subordinate server configuration is called *rollup mode*. When you enable rollup mode on the master server, you can see capacity and status information in the rollup resource pages for the storage systems, servers, hypervisors, switches, and fabrics that are monitored by your rollup servers.

You can also configure a master rollup server as a subordinate server to other master servers. This allows you to create a comprehensive rollup reporting structure across your enterprise.

Note: If you configure a master rollup server as a subordinate of another master server, the data collected by that master server is not rolled up to the higher master server. Master rollup servers only collect and display data from their immediate subordinate servers. For example, if you configure master server 2 as a subordinate to master server 1, the data that is rolled up to master server 2 from its subordinate servers is not collected by master server 1. Master server 1 only collects capacity and status information about master server 2, not the subordinate servers monitored by master server 2.

Note: After you upgrade a master rollup server you must run probes of the subordinate servers that were added to the master server prior to the upgrade. Use the Start Probe action to run probes of the subordinate servers. Wait for the first probe (manual or scheduled) to complete, before you work with a new function.

Managing subordinate servers

After you add a subordinate server on the **Settings** > **Rollup Server Connections** page, you can view information about the server and perform actions on it. You must have IBM Spectrum Control Administrator role authorization to add, remove, and manage subordinate servers on this page.

Information about subordinate servers

IBM Spectrum Control provides the following information about the subordinate servers on the **Rollup Server Connections** page:

Managed resources

The number of block storage systems, servers, switches, fabrics, and hypervisors that the subordinate server is managing. You can view information about these resources on the rollup resource list pages for the resources, such as the **Switches (Rollup)** page. NPV switches are not included in this number on either the master or subordinate Switches (Rollup) resource list pages.

Data collection

The status of the last probe of the subordinate server, and the time of the next probe of the server. Hourly probes are automatically scheduled when a subordinate server is added to a master server.

Actions

You can perform the following actions on the subordinate server. Unless otherwise noted, you must have Administrator role authority to perform these actions.

Open logs

Opens the probe logs window for the subordinate server in a separate browser window. Any IBM Spectrum Control user role can perform this action.

Start probe

Starts a probe of the subordinate server. If a probe is already running for the server, you must wait until the current probe completes before starting a new probe.

Remove server

Removes the subordinate server from the **Rollup Server Connections** page and the **Servers** (**Rollup**) list. Also removes any associated resources that are managed by the subordinate server. For example, if a subordinate server is the managing instance for a switch and the subordinate server is removed from the master server, the switch will no longer display on the **Switches** (**Rollup**) resource list page.

Note: Removing a subordinate server on the **Rollup Server Connections** page does not remove that server from the non-rollup **Servers** page. If the server has been added by using the **Add Server** dialog on the non-rollup Servers page, IBM Spectrum Control will continue to manage that server unless you remove it from the Servers page.

Modify connection

Allows you to change the host name, user name, password, and port of the subordinate server.

Rollup mode

Rollup mode is a special view of your enterprise and is enabled by configuring your IBM Spectrum Control servers to communicate and gather data through server probes. Rollup mode requires a master server configured to collect capacity and status information from one or more subordinate servers. Once you

have configured a master server with at least one subordinate server, you can go to the rollup icon the menu bar to enter rollup mode. The rollup icon does not display until you have configured at least one subordinate server for the master (local) server. You might have to refresh your browser window to see the rollup icon after you add the first subordinate server for the master server you are configuring.

Entering rollup mode

To view rollup capacity and status information from your IBM Spectrum Control servers, you must enter rollup mode.

1. On the **Rollup Server Connections** page, add at least one subordinate server. Subordinate servers must have IBM Spectrum Control installed and running.

2. Move the mouse pointer over the rollup icon in the menu bar and select **Enter rollup mode**.

You will see the following changes in the menu bar:

- The color of the rollup icon changes:
- The lower edge of the menu bar changes to orange.
- The menu bar options change to Storage, Servers, Network, Rollup Server Connections.

Using rollup mode

Rollup mode provides lists of resources that are managed by subordinate servers. For example, you can see a list of all the network switches that are managed by a master server and its associated subordinate servers. These resource list pages are similar to the pages you work with when not in rollup mode. However, in some cases rollup mode provides different data columns and different actions that you can take on the rollup storage devices.

View rollup information

- 1. In the rollup menu bar, select the type of storage device you want to view information about. For example, **Servers**.
- 2. Select one of the available menu options. The Servers menu bar item has two options: **Servers** and **Hypervisors**.
- 3. On the resource list page, you can see the data about that resource type that has been collected from the master and subordinate servers.
- 4. Rollup resource lists function similar to non-rollup pages: you can add and reorder data columns, right-click items in the resource list to perform actions on them, and other familiar functions. You can also use Ctrl-click to select multiple list items.

Actions

You can perform actions on items in the resource lists by right-clicking them or by selecting one or more list items and clicking **Actions**.

View properties

View the properties of the selected item in the resource list. This action is not available for multiple selections at the same time.

Export

Export the data for one or more list items. You can specify PDF, CSV, or HTML output.

Units of measurement for storage data

IBM Spectrum Control uses decimal and binary units of measurement to express the size of storage data.

Decimal units such as kilobyte (KB), megabyte (MB), and gigabyte (GB) are commonly used to express the size of data. Binary units of measurement include kibibyte (KiB), mebibyte (MiB), and gibibyte (GiB). <u>Table</u> 24 on page 79 compares the names, symbols, and values of decimal and binary units.

Table 24. Comparison of binary and decimal units and values						
Binary			Decimal			
Name	Symbol	Value (base 2)	Name	Symbol	Value (base 10)	
kibibyte	KiB	2 ¹⁰	kilobyte	КВ	10 ³	
mebibyte	MiB	2 ²⁰	megabyte	МВ	10 ⁶	
gibibyte	GiB	2 ³⁰	gigabyte	GB	10 ⁹	
tebibyte	TiB	2 ⁴⁰	terabyte	ТВ	10 ¹²	
pebibyte	PiB	2 ⁵⁰	petabyte	РВ	10 ¹⁵	
exbibyte	EiB	2 ⁶⁰	exabyte	EB	10 ¹⁸	

Binary units of measurement express the size of data more accurately. When you compare the size of 100 KB to 100 KiB, the difference is relatively small, 2.35%. However, this difference grows as the size of the data values increases. When you compare the size of 100 TB to 100 TiB, the difference is 9.06%.

In general, IBM Spectrum Control uses base 2 values for memory and disk space values, and base 10 values for space on physical hard drives.

<u>Table 25 on page 80</u> shows the percentage difference between decimal and binary values across a range of data sizes.

Table 25. Percentage difference between decimal and binary units					
Decimal value	Binary equivalent of decimal value	Difference			
100 kilobytes (KB)	97.65 kibibytes (KiB)	2.35%			
100 megabytes (MB)	95.36 mebibytes (MiB)	4.64%			
100 gigabytes (GB)	93.13 gibibytes (GiB)	6.87%			
100 terabytes (TB)	90.94 tebibytes (TiB)	9.06%			
100 petabytes (PB)	88.81 pebibytes (PiB)	11.19%			
100 exabytes (EB)	86.73 exbibytes (EiB)	13.27%			

Role-based authorization

Roles determine the functions that are available to users of IBM Spectrum Control. When a user ID is authenticated to IBM Spectrum Control through the GUI, CLI, or APIs, membership in an operating system or LDAP group determines the authorization level of the user.

The following table shows the IBM Spectrum Control roles and their authorization levels:

Table 26. IBM Spectrum Control roles and authorization levels		
Roles	Authorization level	
Administrator	This role has full access to all monitoring and administrative functions. At least one group must have the Administrator role.	
	Note: When IBM Spectrum Control is first installed, the following operating system groups are assigned the Administrator role:	
	Windows: Administrators	
	UNIX and Linux: root	
	AIX: system	

Table 26. IBM Spectrum Control roles and authorization levels (continued)			
Roles	Authorization level		
Monitor	This role has access to the following read-only functions:		
	Viewing and exporting information about monitored resources		
	Viewing, acknowledging, and removing alerts		
	Viewing tasks and data collection jobs		
	Opening management GUIs		
	Opening logs		
	Viewing chargeback, consumer, predefined capacity and inventory, and custom reports		
	Exception: Users with the Monitor role can provision storage if they are granted permission in a service class. A service class is a logical entity that describes storage capabilities and characteristics and can be used to specify requirements for storage provisioning. For more information about service classes, see .		
External Application	If you assign the External Application role to the user, you must also assign one or more service classes to the user.		
	This role does not enable users to log in to the IBM Spectrum Control GUI.		

Tips:

- ullet To determine the role of the user who is logged in, click the user icon lacksquare In the upper-right corner of any page in the GUI.
- If a user belongs to multiple groups and the groups have different roles, the role with the highest level of authorization is granted to the user. For example, if a user belongs to a group that is assigned the Administrator role and also belongs to a group that is assigned a Monitor role, the user is granted the authorization of the Administrator role.
- If a user is not a member of a group that is assigned a IBM Spectrum Control role, no access is granted to that user.
- If assigned the Monitor role, a user can only open and view logs from the Data Collection page for the selected resource.

Nested groups are not supported: Adding active directory or any other type of domain user group to a local operating system group is not supported in IBM Spectrum Control. You can configure IBM Spectrum Control to authenticate domain IDs that rely on the operating system to perform the authentication operation against the active directory, but it cannot resolve nested groups.

Alternatively, you can configure LDAP authentication to perform queries against active directory user repositories and assign domain groups directly to roles within IBM Spectrum Control.

Fabrics and zones

Use the IBM Spectrum Control to learn more about fabrics and zones.

Related tasks

Provisioning with zone control

When you provision storage, you can set a zoning policy to enable automatic zoning. When automatic zoning is enabled, IBM Spectrum Control can create zones during storage provisioning to connect a server to a storage system.

Monitoring fabrics and zones in the GUI

You can monitor switches and fabrics in the IBM Spectrum Control GUI. You can also enable automatic zoning so that new zones are created during storage provisioning, if necessary, to connect a server to a storage system.

You can use the GUI to add switches and fabrics for monitoring. You can add the following types of switches:

- Brocade
- Cisco

After switches and fabrics are probed, you can view detailed status and performance information about the switches and fabrics. You can also test the connection to the switch or fabric.

Information that you can view about a fabric includes the following properties:

- The fabric type
- The number of physical and virtual switches in the fabric
- The number of ports that are on all the switches in the fabric
- The number of switch ports that are online and connected to other ports in a fabric
- The name of the principal switch of the fabric

Information that you can view about a switch includes the following properties:

- The number of ports on a switch
- The number of ports that are connected to a storage resource
- The name of the fabric where a switch is a member
- The name of the vendor or manufacturer for a switch

From the **Details page** of a fabric, you can open the **Zone Sets page** to view information about the zone sets on the fabric, and to identify which is the active zone set.

You can enable automatic zoning in the GUI. When automatic zoning is enabled, IBM Spectrum Control creates new zones during provisioning if new zones are needed to connect the storage system with the server.

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

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.

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 FIBM Spectrum Control interoperability matrix for switches.

Zoning for Cisco MDS 9000 switches

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. IBM Spectrum Control provides basic zone discovery for the Cisco MDS 9000 series switches.

IBM Spectrum Control supports these zone member types:

- N_Port WWN
- FC ID (the fibre channel ID of an N_port attached to the switch)

IBM Spectrum Control does not allow zone management of zones that contain unsupported members.

Brocade switches in zones

Brocade switches that are managed by using Brocade Network Advisor 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, IBM Spectrum Control preserves them and does not modify them in any way.

You can create, change, and delete non-standard zones by using the Brocade switch management application.

Note: IBM Spectrum Control 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 Spectrum Control uses agents to gather data: Common Information Model (CIM) agents, Storage Resource agents, and SNMP agents.

SMI-S providers

SMI-S providers enable communication between storage devices, such as storage systems and fabric switches, and IBM Spectrum Control. Each storage device that you want to manage must have an SMI-S provider either installed or embedded unless it is a device that uses Native API connection.

SMI-S providers are provided by the vendor of the storage subsystem or fabric switch. 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 SMI-S provider enables communication between the storage device and IBM Spectrum Control. Commands and responses are transmitted between IBM Spectrum Control and the SMI-S provider using an XML transport layer. The SMI-S provider to storage device layer uses a vendor-specific proprietary interface.

The SMI-S provider usually must be installed and configured, so that it can identify the storage devices with which it communicates. Some storage devices, such as fabric switches, contain embedded SMI-S providers and so do not require that SMI-S providers be installed. In these cases, IBM Spectrum Control must be configured to point directly to the storage devices that contain the embedded SMI-S providers.

SMI-S providers can be referred to as CIM agents, CIM proxy agents, and CIMOM agents. SMI-S providers can be embedded in the device or installed on a separate computer.

Note:

- Do not install multiple SMI-S providers on a single computer because of port conflicts.
- Do not install a SMI-S provider on the system where a IBM Spectrum Control server component is installed.

Storage Resource agents

Use Storage Resource agents to collect information about storage resources such as servers, virtual machines, workstations, and HBAs.

You must deploy Storage Resource agents on resources where you want to gather the following information:

- · Asset information
- File and file system attributes
- Network-attached storage (NAS) device information

Tip: IBM Spectrum Control creates and updates agentless servers automatically after it probes storage systems and hypervisors.

SNMP agents

An SNMP agent, also called an out-of-band agent, is software that runs on a switch and can be used to monitor and manage systems and devices in your network. IBM Spectrum Control uses SNMP to retrieve

information about specific switches, such as Cisco switches, and the fabrics that the switches are members of.

IBM Spectrum Control uses SNMPv3 (preferred) or SNMPv1 to probe Cisco switches to collect performance data. Some switches are configured to use SNMPv3 by default.

IBM Spectrum Control REST API

You can use the Representational State Transfer (REST) API for IBM Spectrum Control to access information about resources and to generate custom capacity, configuration, and performance reports.

To get the information that you need about your resources, you can use a REST command line utility or you can use a web browser.

Chapter 5. Tutorials

Use tutorials to learn how to perform specific tasks and become familiar with IBM Spectrum Control.

The provided tutorials set out situations that can be solved by using IBM Spectrum Control. When using these tutorials you should remember the following:

- The tutorials are intended to be step-by-step instructions that you follow to complete a specific job. Each tutorial contains multiple tasks that must be followed in the order given.
- Specific values are given in the tutorials, 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 tutorials do not provide in-depth conceptual information about the tasks. See the appropriate topics in the IBM Spectrum Control Knowledge Center for more information about the concepts behind the tasks.
- You will complete a tutorial successfully if you follow the instructions correctly and your system performs as expected.

Optimizing the performance of storage virtualizers

In this tutorial, you balance pools and analyze tiering to help resolve performance hotspots and manage storage resources efficiently.

You can optimize the placement of volumes on the following storage virtualizers:

- SAN Volume Controller
- Storwize V7000
- Storwize V7000 Unified block storage

Learn more: For information about the components that can be used with IBM Spectrum Control Version 5.2 (or later), see

IBM Spectrum Control interoperability matrix at http://www.ibm.com/support/docview.wss?uid=swg21386446.

To optimize the placement of volumes in storage virtualizer pools, you must complete the following tasks:

- Probe the storage virtualizers.
- Set the tier level of the storage pools that you want to analyze.
- Collect performance data to analyze the pools and volumes.

When you run the analysis to balance pools, you optimize storage performance by redistributing volume workloads across pools on the same tier.

When you run the tiering analysis, you optimize storage performance by specifying thresholds to move volumes to higher and lower tiers.

Related tasks

"Optimizing storage tiering" on page 438

To optimize the placement of volumes on storage tiers, analyze the tiering of volumes in your storage environment.

"Optimizing storage pools" on page 444

Analyze the activity of pools and resolve performance hot spots by redistributing volumes across each storage tier.


"Managing tasks for tiering storage, balancing pools, and transforming storage" on page 461
To optimize the resources in your storage environment, use tiering storage, balancing pools, and transforming-storage tasks. Use the **Tasks** page to manage all the tasks that are used to optimize storage.

Use resource list and resource detail pages to manage the optimization tasks for resource types and for specific resources.

Tutorial: Analyzing and re-tiering volumes in pools on tier 1

In this tutorial task, you analyze the current placement of volumes in pools on tier 1 so that you can move volumes with low workloads from tier 1 to tier 2 or tier 3 pools.

The pools that you want to analyze are on a SAN Volume Controller with two DS8000 storage systems. The pools on the storage virtualizer are assigned to tier 1, tier 2, and tier 3.



For example, you want to analyze the data that is collected over 10 days and when the volumes are most active. So, you set the period for the analysis to 10 days and exclude the data that is collected at the weekend. You also exclude the data that is collected before 9 AM and after 6 PM because you know that the volumes are most active from 9 AM to 6 PM.

Tiering and volume placement criteria 1 Tiering criteria The source volumes that are selected are in pools on Tier 1, and the target pools that are selected are on Tier 1, Tier 2, and Tier 3. The tiering thresholds for the volumes are >1000 I/O per second for Tier 1 and >500 I/O per second for Tier 2. **Tiering Criteria Tier 1: Volume Analysis** Tiering recommendations Tier 1 I/O Rate: >1000 Place volume in pool on Tier 2 I/O Rate = 900 Tier 2 I/O Rate: <1000 and >500 Tier 3 I/O Rate: <500 Place volume in pool on Tier 3 2 Volume placement criteria To ensure that the pools that are selected as target pools can handle the additional workload, you specify the maximum I/O rates for the target pools on each tier. **Below Threshold Add Volume Target Pools** Target Pool 1 Target Pool 2

You want volumes with the highest workloads in tier 1 pools and you want volumes with low workloads in tier 2 or tier 3 pools. To tier your storage, you enter two I/O rate thresholds to generate recommendations to re-tier the volumes.

Tip: You can tier volumes by I/O rate or I/O density.

To ensure that the pools that are selected as destination pools for the volumes can handle the additional workload, you specify the maximum I/O rates for pools on each tier of storage. Only the pools with I/O rates below the maximum I/O rate can be selected as destination pools for the volumes.

- 1. From the Storage menu, click Storage Systems.
- 2. Right-click the storage virtualizer and click View Details.
- In the navigation pane, click Pools.
 To organize the pools by tier level, click the Tier column heading.
- 4. Select the pools on tier 1, right-click and then click Analyze Tiering.
- 5. Select all of the pools on tier 2 and tier 3 as target pools, and then click Next.
- 6. Select the period that you want to use to analyze the performance data.

- 7. Set the threshold for moving volumes from tier 1 to tier 2.
 - The volumes with I/O rates above the threshold that you set remain in the tier 1 pools. The volumes with I/O rates below the threshold that you set are moved down to tier 2 pools.
- 8. Set the threshold for moving volumes from tier 2 to tier 3.
 - The volumes with I/O rates above the threshold that you set remain in the tier 2 pools. The volumes with I/O rates below the threshold that you set are moved down to tier 3 pools.
- 9. Set the maximum I/O rate for pools on each tier.
- 10. Click Analyze.

The **Analyze Tiering** task is shown.

Click **Execute** to implement the recommendations. To create a schedule, click **Schedule** > **Analysis**. Alternatively, you can modify the criteria that you set for analyzing the volumes on the **Tasks** page and run the analysis again.

Related tasks

Modifying the criteria for analyzing tiering

When a task is created after you run the tiering analysis, you can modify the criteria for tiering the volumes.

Tutorial: Collocating volumes

In this tutorial task, you want to minimize the exposure of servers to multiple back-end storage systems by collocating volumes that are assigned to the same hypervisor or server. You can enforce the collocation of volumes when you enter the criteria for analyzing tiering and balancing pools.

To ensure that volumes in the same storage pool that are assigned to the same server or hypervisor are kept together, you want to enforce the collocation of volumes. By enforcing the collocation of volumes, you prevent the placement of related volumes in destination pools that might be on multiple back-end storage systems.

Multiple host connections to the same hypervisor or server: If the volumes in the source pool that are assigned to the same hypervisor or server are assigned to different host connections, the collocation of the volumes is affected. In such cases, if volumes require optimization, the volumes that are assigned to the same host connection are kept together. To view information about the host connection for the volume, right-click the volume, select **View Details**, and then click the **Host Connections** tab.

- 1. From the **Storage** menu, click **Pools**.
- 2. Right-click the pools that you want to analyze, and then click Analyze Tiering.
- 3. Select the target storage pools.
- 4. On the Optimize the Placement of Volumes page, ensure that Collocate volumes is set to Yes.
- 5. Click Analyze.

If a volume requires re-tiering, a recommendation is generated to move all of the volumes that are assigned to the same server to the same destination pool. If none of the target pools have sufficient space to accommodate all of the volumes, then recommendations to move the volumes are not generated.

Related tasks

Optimizing storage pools

Analyze the activity of pools and resolve performance hot spots by redistributing volumes across each storage tier.

Optimizing storage tiering

To optimize the placement of volumes on storage tiers, analyze the tiering of volumes in your storage environment.

Modifying the criteria for analyzing tiering

When a task is created after you run the tiering analysis, you can modify the criteria for tiering the volumes.

Monitoring capacity usage at different levels of a business hierarchy

The monitoring and management of applications and departments enables you to monitor storage capacity usage, recognize trends, monitor health status, and troubleshoot performance of the storage resources in your business organization.

The department model comprises these main elements:

- The department
- The subdepartments that the department contains
- The applications, or the application subcomponents, or both that the department uses

In the following tutorial, you want to create a department model that monitors the capacity and space usage of the Books Sales department, and these subdepartments:

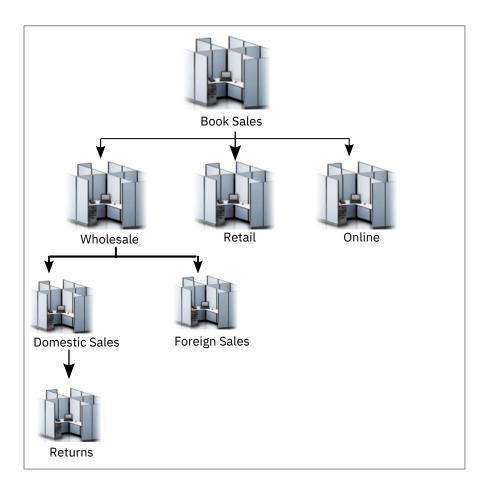
- Wholesale
- Retail
- Online

To monitor the capacity and space usage of the departments, you add the applications that the department and its subdepartment use to each of the subdepartments.

The department model that you create determines how you can view capacity usage and space information. To view the total capacity and space usage of a department regardless of the subdepartments that the department contains, you create a department and add the applications and subcomponents to the department.

If, however, you want to see the capacity and space usage of the department and its subdepartments, you create a department hierarchy by adding subdepartments to departments. You then add the applications and application subcomponents that each subdepartment uses to the subdepartments. You can then view capacity and space usage for the department and for the individual subdepartments.

You can create more complex models of departments by adding subdepartments to subdepartments. In the following illustration, the department model is extended to include two more layers of subdepartments under the Wholesale subdepartment.



To monitor the capacity and space usage of the departments in your business organization, complete the following tasks:

- Create the Book Sales department.
- Create the Wholesale, Retail, and Online departments to represent the subdepartments.
- Add the Wholesale, Retail, and Online as subdepartments to the Book Sales department.
- Create the Book Sales DB application and the Wholesale Transactions, Retail Transactions and Online Transactions applications to represent the subcomponents that are used by the departments.

Different applications require different storage capabilities, and different levels of performance and uptime. The storage requirements of the departments are constantly growing and they need to modify their behaviors in a way that justifies the cost of their actions concerning storage utilization. The data modeling of storage resources using the application and department concept enables you to plan and implement a chargeback system if necessary.

Related concepts

"Departments" on page 406

View information about departments and structure the hierarchy to mirror your business organization. Monitor the detailed information about capacity and space usage that is collected for departments and subdepartments, and monitor the performance of the applications and the application subcomponents that are added to departments and subdepartments.

Related tasks

"Creating departments" on page 407

To model the storage capacity that is consumed in a department for your business environment, in accordance with other department and application members, create departments for monitoring in IBM Spectrum Control.

"Adding and creating subdepartments" on page 408

To further enhance your data modeling at the department level and to add to your overall business hierarchy for monitoring, add an existing subdepartment or create a new subdepartment.

"Viewing information about departments" on page 410

You can view detailed information about departments for use in storage capacity monitoring and performance troubleshooting. You can view additional information on the applications and subdepartments that are used to establish a hierarchal business structure.

"Creating applications" on page 400

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies for applications, you can create applications for monitoring in IBM Spectrum Control.

"Adding subcomponents" on page 403

You created a new application and a subcomponent, with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

Tutorial: Comparing storage usage in each department

In this tutorial, you compare the storage usage in each department within your organization to spot a potential storage usage issue.

To monitor the capacity and space usage of a department, you create the department and subdepartments, and you add the application subcomponents to the subdepartments.

You want to monitor the capacity and space usage of departments in an organization that sells books.

The sales department of the organization, Book Sales, has these subdepartments:

- Wholesale
- Retail
- Online

Although you can add applications and application subcomponents when you create departments, it is easier and quicker to create the applications and application subcomponents beforehand.

You want to know the total amount of storage space that the Book Sales department uses, and you want to know the amount of space that each subdepartment uses. To know how much space the Book Sales and its subdepartments use, you must associate the department and its subdepartments with the applications that they use. For example, the Book Sales department and its subdepartments use a database application named Book Sales DB. The Book Sales DB application contains these subcomponents:

- · Wholesale Transactions
- Retail Transactions
- Online Transactions

To create the department and subdepartments, you complete these tasks:

- 1. You create the Wholesale department and add the applications and application subcomponents that the subdepartment uses.
- 2. You create the Retail department and add the applications and application subcomponents that the subdepartment uses.
- 3. You create the Online department and add the applications and application subcomponents that the subdepartment uses.
- 4. You create the Book Sales department.

You do not have to add applications or subcomponents to the Book Sales department because the Book Sales department inherits the applications and application subcomponents that are added to the subdepartments.

5. You add the Wholesale, Retail, and Online departments as subdepartments to the Book Sales department.

When you complete adding the departments and associating the departments with the applications that are used by the departments, you can view the information that is collected about the departments on the Departments page and on the details page for the department.

You want to target what department and what storage resources might be responsible for a change or a upward trend in significant storage usage. By correctly classifying the department or departments you can properly plan to archive or switch the storage within the organization.

Related concepts

"Departments" on page 406

View information about departments and structure the hierarchy to mirror your business organization. Monitor the detailed information about capacity and space usage that is collected for departments and subdepartments, and monitor the performance of the applications and the application subcomponents that are added to departments and subdepartments.

Related tasks

"Creating departments" on page 407

To model the storage capacity that is consumed in a department for your business environment, in accordance with other department and application members, create departments for monitoring in IBM Spectrum Control.

"Adding and creating subdepartments" on page 408

To further enhance your data modeling at the department level and to add to your overall business hierarchy for monitoring, add an existing subdepartment or create a new subdepartment.

"Viewing information about departments" on page 410

You can view detailed information about departments for use in storage capacity monitoring and performance troubleshooting. You can view additional information on the applications and subdepartments that are used to establish a hierarchal business structure.

"Creating applications" on page 400

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies for applications, you can create applications for monitoring in IBM Spectrum Control.

"Adding subcomponents" on page 403

You created a new application and a subcomponent, with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

Using applications and subcomponents to monitor capacity and space usage

To monitor the performance, capacity, and space usage of the applications in your business organization, create applications and subcomponents. You can also add applications to departments so that storage capacity and usage can be monitored in an overall business hierarchal manner.

You can create a simple application model, which consists of an application and use a filter to associate the storage resources that the application uses with the application. Alternatively, you can create a complex application model that comprises an application and application subcomponents and then add filters to associate the storage resources that each application subcomponent uses with each of the subcomponents.

Besides creating applications to monitor the performance, capacity, and space usage of the applications in your business organization, you also create applications to add them to departments. When you add applications to departments, you can monitor the capacity and space usage of the departments and you can monitor the performance of the storage resources that are associated with the applications that each department uses.

To monitor the capacity of an application, you create an application model that is based on the following elements:

The application

- The subcomponents that the application contains
- The applications, application subcomponents, or both that are added to the department

In the following tutorial, you want to create an application model that monitors the capacity, space usage, and performance of the Book Sales DB application. The application model that you want to create comprises these application subcomponents:

- Wholesale Transactions
- Retail Transactions
- · Online Transactions




Figure 1. An application with a three subcomponent hierarchy

To understand how capacity is trending for storage in an application and subcomponents, complete these tasks:

- 1. Create the Book Sales DB application.
- 2. Create the Wholesale Transactions, Retail Transactions and Online Transactions applications to represent the subcomponents.
- 3. Add the Wholesale Transactions, Retail Transactions, and Online Transactions as subcomponents to the Book Sales DB application.
- 4. Create the Book Sales department and the Wholesale, Retail, and Online subdepartments to be associated with the applications.

Related concepts

"Applications" on page 398

View information about storage area network resources that you can model by using applications in your business environment. The monitoring and management of applications enables you to perform capacity trending, health monitoring, and performance troubleshooting tasks.

Related tasks

"Creating applications" on page 400

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies for applications, you can create applications for monitoring in IBM Spectrum Control.

"Adding subcomponents" on page 403

You created a new application and a subcomponent, with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

"Viewing information about applications" on page 404

You can view information about the application that you created as related resources, filters that are used to populate the application with resources, any subcomponents, capacity trending and performance information.

Tutorial: Viewing storage capacity and usage trends

You can view the capacity and usage trends of storage resources used by an application and subcomponents to diagnose how much storage a particular application is using in relation to the other applications and departments in the business hierarchy.

You want to understand how capacity is trending for storage that is used by the Book Sales DB application and the Wholesale Transaction, Retail Transactions and Online Transactions subcomponents and the association the application has with the Book Sales department and subdepartments.

The storage resources that are used by the application and that are used by each subcomponent are associated with the application and its subcomponents. For example, the storage resources are assigned as follows:

- 1. Create a database application called Book Sales DB and assign resources that use the volumes on the IBM® Storwize V7000 with names that begin with **bksales** for storing sales transactions.
- 2. Create the subcomponent applications named Wholesale Transaction, Retail Transactions and Online Transactions.
- 3. Assign the volumes on the IBM Storwize V7000 with names that begin with **bksales_ws** to the Wholesale Transactions subcomponent, which is used for storing wholesale sales transactions.
- 4. Assign the volumes on the IBM Storwize V7000 with names that begin with **bksales_rt** to the Retail Transactions subcomponent, which is used for storing retail sales transactions.
- 5. Assign the volumes on the IBM Storwize V7000 with names that begin with **bksales_ol** to the Online Transactions subcomponent, which is used for storing online sales transactions.

Note: On the Create Filter page, you can specify a name pattern to determine which volumes to include. For volumes, you can specify name patterns to determine from which servers, storage systems, or pools the volumes are selected. You can then click **Preview** to view the volumes that are selected for inclusion in your application.

To monitor space usage for the Book Sales department and its subdepartments, you associate the Book Sales DB application and the Wholesale, Retail, and Online Transactions subcomponents of the application as follows:

- You associate the Book Sales DB application with the Book Sales department.
- You associate the Wholesale subdepartment with the Wholesale Transactions subcomponent.
- You associate the Retail subdepartment with the Retail Transactions subcomponent.
- You associate the Online subdepartment with the Online Transactions subcomponent.



Figure 2. The application, subcomponent, assigned resources and department association

You can view the capacity information that is collected about the application on the Applications page. You can also view capacity information for each application and subcomponent on the details page. When you complete adding the departments and associating the departments with the applications that are used by the departments, you can view the information that is collected about the departments on the Departments page and on the details page for the department.

Related concepts

"Applications" on page 398

View information about storage area network resources that you can model by using applications in your business environment. The monitoring and management of applications enables you to perform capacity trending, health monitoring, and performance troubleshooting tasks.

Related tasks

"Creating applications" on page 400

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies for applications, you can create applications for monitoring in IBM Spectrum Control.

"Adding subcomponents" on page 403

You created a new application and a subcomponent, with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

"Adding applications to departments" on page 409

To further enhance the data storage modeling in a business environment and see the storage capacity that is used at the different levels of your business hierarchy, add applications to departments. The departments you select from the **Add to Department** page become parents of the selected applications.

"Viewing information about applications" on page 404

You can view information about the application that you created as related resources, filters that are used to populate the application with resources, any subcomponents, capacity trending and performance information.

Tutorial: Viewing NPIV connections between server ports and switch ports in a fabric

In this tutorial, a storage administrator views the N-Port ID Virtualization (NPIV) connections between the server ports and the switch port in the fabric. With NPIV, multiple node ports can be logically connected to one switch port.

You want to view connectivity information for a switch port that has NPIV connections.

1. In the menu bar, go to **Network > Switches**. Double-click the switch that you want to view, or right-click the switch and click **View Details**.

Tip: Alternatively, in the menu bar, go to **Network** > **Fabrics**. Double-click the fabric that you want to view, or right-click the fabric and click **View Details**. In the **Internal Resources** section, click **Switches**. Double-click the switch that you want to view, or right-click the switch and click **View Details**.

- 2. In the **Internal Resources** section, click **Ports**. You can view details about the ports, including the number of NPIV connections.
- 3. In the **Connected NPIV Ports** column for the port that you want to view, click the number of ports. You can view the following information about the NPIV connections:
 - In the **Connected NPIV Ports** tab, you can view details about the NPIV connections, such as the port type and the name of the connected resource. The port type can be N_Port (end node port) or NP_Port (proxy node port).
 - In the **Connectivity** tab, you can view details about the physically connected end-node port.

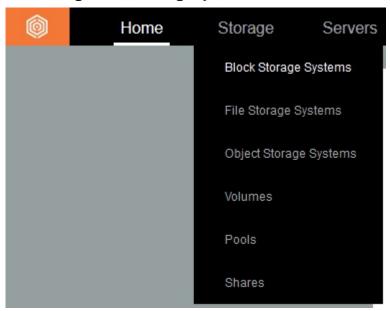
Tip: If IBM Spectrum Control does not differentiate between the physical and logical connections for the switch port, the details about the NPIV connections are displayed in the **Connectivity** tab. The **Connected NPIV Ports** tab is not displayed.

Tutorial: Exporting performance data for a SAN Volume Controller system

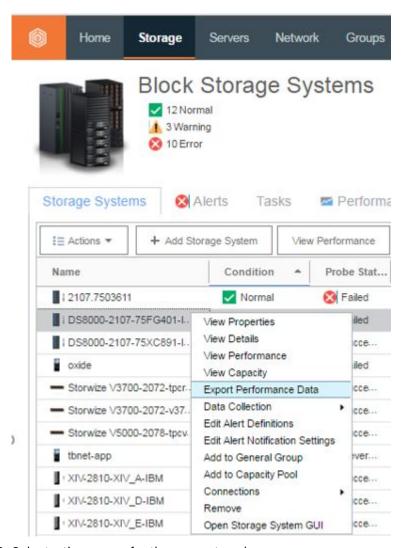
In this tutorial, you export performance data for a SAN Volume Controller to a compressed package. You then send the package to IBM Support.

Some volumes on a SAN Volume Controller system have performance problems. You consult IBM Support who require detailed performance data about the SAN Volume Controller system to diagnose the problem. You are asked to export performance data by using IBM Spectrum Control.

1. Go to Storage > Block Storage Systems.



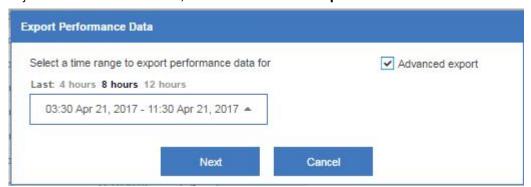
2. Right-click the SAN Volume Controller system, then click **Export Performance Data**.



3. Select a time range for the support package.

Typically, the time range includes the time when the performance problem occurred.

4. If you are instructed to do so, select the **Advanced export** check box. Click **Next**.



5. Click Create when you are prompted to create the package with the details that you specified.

The package is created and your browser downloads the package.

Related reference

Exporting performance data for storage systems and fabrics


To help resolve performance issues with storage systems and fabrics, you can export performance data for the resources to a compressed file. If you contact IBM Support to help you analyze the problem, you might be asked to send this file.

Tutorial: Comparing the performance of storage systems


In this tutorial, you compare the performance of two storage systems.

Your company recently purchased an IBM FlashSystem 840 storage system. The storage system was added to IBM Spectrum Control and the performance of the storage system is being monitored. Your managers ask you to compare the performance of the FlashSystem 840 storage system to another storage system so that they can assess the performance of their investment. In particular, your managers are interested in the response times and the I/O rates of the FlashSystem 840.

1. Go to Storage > Block Storage Systems.



- 2. Click the FlashSystem 840 storage system and the non-flash storage system that you want to compare the FlashSystem 840 with.
- 3. Right-click the selected storage systems, then click **View Performance**.



A performance view is displayed. The view shows the performance of the two storage systems on a chart and in a table. By default, the performance view shows the Total I/O Rate - overall and Overall Response Time metrics. The time period for the information is the last 12 hours.



Tutorial: Reviewing and updating your agentless servers

In this tutorial, you review the agentless servers that are created automatically by IBM Spectrum Control, and update your agentless servers.

Every month, you create a chargeback report for the physical servers in your storage environment. The report shows how much capacity is used by physical servers, and the cost of that capacity.


IBM Spectrum Control creates and updates agentless servers automatically after storage systems and hypervisors are probed. The agentless server information is used to create an accurate model of your storage environment, which is then used by the chargeback report to show the consumption of storage by server. You want to review your agentless servers to make sure that they accurately represent your storage environment.

- 1. Go to Servers > Servers.
- 2. Review the servers that have **Agentless** in the **Condition** column.

 You notice that IBM Spectrum Control identified 4 host connections and created an agentless server for each host connection. However, in your environment, these 4 host connections are on one server. You need to combine the 4 agentless servers into a server.
- 3. Select the four agentless servers, right-click, and then click **Merge Servers**.








- 4. Type a name for the new server, then click Merge.
- 5. Review the agentless servers again.

You notice that IBM Spectrum Control created an agentless server from a group of host connections that appear to be related. However, in your environment, these host connections are actually individual servers. You need to separate the agentless server into individual servers.

6. Select the agentless server, right-click, and then click **Separate Server**.





- 7. If any WWPNs on the server are used by more than one host connection, an error message is displayed. The message displays the WWPNs that are used by more than one host connection. Review the WWPNs and make sure that each WWPN is used by only one host connection. Then, try to separate the server again.
- 8. On the **Separate Servers** dialog, click **Apply**.

 On the resource list page for servers, an agentless server is created for each of the servers.

Related concepts

Agentless servers

IBM Spectrum Control creates and updates agentless servers automatically after it probes storage systems and hypervisors.

Tutorial: Troubleshooting performance

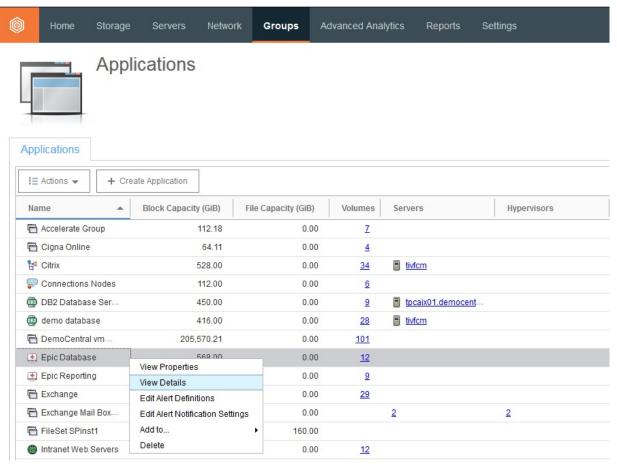
Bob is a storage administrator. One morning, he receives a ticket: A critical medical application, **Epic Database**, has a performance problem. Bob must investigate the problem and identify the cause.

Previously, Bob created an **Epic Database** application on IBM Spectrum Control and specified the resources that the application uses.


1. Bob clicks **Applications** on the dashboard to troubleshoot only the **Epic Database** resources.



2. Bob right-clicks the Epic Database application on the Applications page, then clicks View Details.



Bob immediately notices the high response times of one volume, vsc-epic-db-test-goldp1, in the **Most Active Volumes** chart. He sees at a glance that this volume is well above 10 ms/op.



Bob is familiar with the response times and other network conditions that are acceptable for the application. He knows that this response time is too high for this application. He also notes that the high response time is sustained, not just spikes. Bob decides to look more closely at the performance of the volume.

- 3. Bob clicks Volumes in the Related Resources section, then clicks the Performance tab.
- 4. Bob uses the chart and table in the **Performance** tab to further investigate the cause of the problem with the **Epic Database** application.

He uses the chart and table in the **Performance** tab to investigate by using the following techniques:

- Change which metrics are displayed in the performance charts. He clicks a metric in the chart legend to display only that metric. He also clicks the **Select Chart Metrics** icon to add other metrics to the chart.
- Change which resources are displayed in the performance charts. He clicks a resource in the legend to display only that metric. He also clicks resources in the table to add them to the chart.
- View the relationships between resources.
- Compare metrics by time range. To make this comparison, he can select either a predefined time range or a custom time range.
- Compare different metrics within the same time range.
- Compare the same metrics across different time ranges.
- Identify the best or worst values for a performance metric from a large list of resources. For example, Bob might want to see which of the managed disks is handling the most workload, or which volumes have the highest response times.
- Compare performance charts side by side. To compare charts, Bob opens a chart in a new window and arranges the charts beside one another.
- Synchronize all charts so that they all show data for the same time period.
- Export the data from a chart in CSV format.

To see how Bob investigates the performance problem, watch the <u>IBM Spectrum Control Performance</u> troubleshooting techniques (6:49) video on YouTube at https://youtu.be/4J1jh5hYV-U.

Tutorial: Monitoring IBM Spectrum Scale performance

You can use IBM Spectrum Control to monitor the performance of the IBM Spectrum Scale clusters in your storage environment.

In the following video, you learn how to use IBM Spectrum Control to compare the performance of multiple IBM Spectrum Scale clusters and then analyze the clusters that have the heaviest I/O loads or the longest response times.

Watch the Monitoring IBM Spectrum Scale performance with IBM Spectrum Control (5:38) video on YouTube at https://youtu.be/9kPbDo9bQSg.

In the video, you learn how to use the following techniques to identify performance issues and bottlenecks in your storage:

- Analyze the performance of all your clusters over a period to determine which clusters have the heaviest I/O loads or the longest response times.
- View which of the nodes that mount a file system have the heaviest I/O loads and highest CPU utilization.
- Compare the file systems on a cluster to determine which file systems have the heaviest I/O loads and the longest response times.
- Monitor the performance of SAN attached volumes that are backing the Network Shared Disks (NSDs) in the cluster. If you have block storage devices, such as IBM Spectrum Accelerate or IBM FlashSystem, that are providing storage for your clusters, you can monitor the block volumes that are backing the NSDs.

Tutorial: Viewing the aggregated workload for an application

Bob is a storage administrator. Bob wants to easily evaluate the performance of the accounting application because the Accounts department is expanding to a new location and they want to replicate their data across both locations. Bob needs to work out the inter-site link capacity requirements.

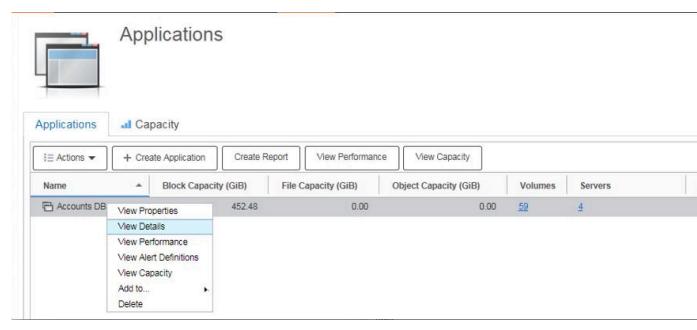
Previously, Bob created an accounting application, **Accounts Database**, and specified the resources that the application uses.

1. Bob clicks **Applications** on the dashboard to view only the **Accounts Database** resources.






2. Bob right-clicks the **Accounts Db** application on the **Applications** page, then clicks **View Details**.




Bob can use the **Aggregated Workload** chart to quickly evaluate the throughput required for the link between both sites.

Overview



Bob can toggle between seeing the aggregated I/O rate and the data rate for the application using the arrow ③ ② icons either side of the chart.

Overview



3. The rates are aggregated on the primary volumes associated with the application which allows Bob to accurately assess the required MiB/s throughput he will need between both locations.

Tutorial: Identifying the source of slow drain problems caused by depletion of buffer credits

Use this tutorial to find out how to use IBM Spectrum Control to identify a host that has depleted buffer credits that are causing a slow drain condition.


Fibre Channel (FC) networks use buffer credits to control the flow of data frames from port to port. The number of buffer credits for a port is the number of data frames that the port can receive. When that number is reached, ports cannot send further data frames until the receiving port indicates that it is ready. If all of the buffer credits of a port are being used, then the port cannot receive more data.

For example, if a host has a performance problem, then its ports might not be able to clear their buffer credits to receive more data. If the host ports cannot receive data, then switch ports cannot send data to the host ports, so the buffer credits of the switch port become depleted too. Ports on other switches in the fabric that try to send data through the switch port are also affected, and their buffer credits become depleted in turn. In this way, the buffer credit problem builds throughout the storage environment. The buffer credit depletion on the host ports impacts the switches that communicated with the host. The switches that communicate with that switch cannot use their buffer credits, so storage systems cannot communicate with the switches.


In this way, a single host with a performance problem can impact all the hosts that use the same switches and inter-switch links. This condition is called *slow drain*. Slow drain in your storage environment can manifest as a problem with storage systems rather than with a host.

1. Configure a performance alert for Port Send Delay Time, Port Send Delay I/O Percentage, or Zero Buffer Credit Timer, depending on the storage system.


For example, configure a Port Send Delay Time alert for FlashSystem 9100 that is triggered when the delay is greater than 10 ms/op.



- 2. To view the alerts, click **Home** and then click **Alerts**. If the Port Send Delay Time alert was triggered, note the time of the alert.
- 3. To view information about the affected storage system, click the link in the Resource column.
- 4. In the Internal Resources section of the storage system details page, click Volumes.
- 5. Click the **Performance** tab.
- 6. Set a time period for the performance chart. Set the start time to before the alert occurred and the end time to after the alert occurred.
- 7. Set the chart to display the following metrics:
 - · Read Data Rate
 - · Overall Response time
- 8. Sort the performance table by the Max Total I/O Rate column.
- 9. Click the volume with the highest Total I/O Rate to show the volume in the chart. Verify that the Read Data Rate spiked at the time that the alert occurred.



- 10. Right-click the volume in the performance table, then click **Host Connection Performance**.
- 11. Set the chart to display the following metrics:
 - · Read Data Rate
 - · Overall Response time
- 12. Verify that the spike for the Read Data Rate and the Overall Response Time occurred when the alert was triggered. If the spike occurred at the same time, the host P9_1par5_epictest01 is the source of the slow drain that is causing problems on the storage system.



Tip: If the buffer credit problem occurs in a cluster, there might be multiple hosts that are mapped to the volume. In this case, you must investigate the hosts individually to determine which ones cause the delay in sending from the port.

Chapter 6. Managing resources

Use IBM Spectrum Control to monitor and manage the resources and storage infrastructure in an enterprise environment. In the GUI, you can view the condition, capacity, performance, and relationships of storage resources.

Resources that you can monitor

You can use IBM Spectrum Control to monitor storage systems, servers, hypervisors, fabrics, and switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

Information about object storage systems includes information about their object resources and related resources. Object resources are resources in an IBM Spectrum Scale cluster, such as accounts and containers, that enable the cluster to be used as an object storage system.

Fabrics

A *fabric* is a network of hubs, switches, adapter endpoints, and connecting cables that support a communication protocol between devices. The following table shows the internal resources and related resources that you can monitor for fabrics.

Table 27. Fabric resources that you can monitor			
Fabric icon	Internal resources	Related resources	*Connected resources (to ports)
Physical fabric Virtual fabric	Logical switches Physical switches Inter-switch connections Blades Ports Zone sets	Hypervisors Servers Storage systems NPV switches (Cisco switches in NPV mode and Brocade switches in Access Gateway mode) Discovered ports	All monitored resources that are connected to ports, plus the following resources: Unmanaged resource. The name is determined from a fabric zone alias. Unmanaged resource. The name is determined from a storage system's mapping to a host. Unmanaged resource. The name is determined from one of a storage system's mappings to a host when multiple mappings are found.

*Connected resources (to ports): The icons for any resources that are managed by IBM Spectrum Control and connected to a port are shown in this column. Icons are also shown for connected resources whose names are known from a fabric zone alias or a storage system's mapping to a host (host connection) but are not currently being managed. If both a fabric zone alias and a storage system's mapping to a host are found for the same unmanaged resource, the fabric zone alias is shown as the name.

Hypervisors

A *hypervisor* is a layer of software or a physical device that manages multiple, virtual instances of operating systems on the same hardware. Hypervisors that you can monitor include VMWare vSphere data sources such as ESX, ESXi, and vCenter servers. If a hypervisor is part of a cluster, information about that cluster is also available.

The following table shows the internal resources and related resources that you can monitor for hypervisors.

Table 28. Hypervisor resources that you can monitor			
Hypervisor icon	Internal resources	Related resources	
OS OS OS	Controllers Disks Data stores Virtual machine that is associated with an agentless server Virtual machine that is associated with a Storage Resource agent Virtual machine that was discovered, but isn't associated with an agentless server or Storage Resource agent	Storage systems Fabrics Switches Cluster	

Servers

A *server* is a computer or host that is connected to a network and provides an access point to that network. If a server is part of a cluster, information about that cluster is also available. You can monitor the following internal resources and related resources for servers.

Table 29. Server resources that you can monitor			
Server icon	Internal resources	Related resources	
	Controllers Disks Volume groups File systems and logical volumes Shares and exports	Storage systems Fabrics Switches	

Storage systems

A *storage system* is a device, such as a Redundant Array of Independent Disks (RAID) controller that creates and manages other storage devices. A storage virtualizer is a storage system that virtualizes storage space from internal storage or from another storage system.

Storage systems can be configured for block storage, file storage, a combination of block and file storage, or a combination of file and object storage. The method that a storage system uses for managing data determines the internal resources or object resources that are monitored.

The following tables show the internal resources, object resources, and related resources that you can monitor for block, file, and object storage systems.

Block storage systems

Table 30. Block storage system resources that you can monitor			
Block storage system icon	Internal resources	Related resources	
	Active quorum Disks Drives Enclosures External disks FC ports Host connections I/O groups IP ports Managed disks Modules Modules Pools RAID arrays	Back-end storage systems Fabrics Hypervisors Servers Switches Virtualizer storage systems GPFS clusters	

File storage systems

Table 31. File storage system resources that you can monitor			
File storage system icon	Internal resources	Related resources	
	₩ Disk Controllers	Block storage systems	
	Clusters		
	Disks	GPFS clusters	
	Shares	Object storage systems	
	Filesets	Hypervisors	
	File systems	Servers	
	Network shared disks	Switches	
	Nodes		
	Pools		
	Snapshots		
	Quotas		

Object storage systems

Table 32. Object storage system resources that you can monitor			
Object storage system icons	Object internal resources	Object resources	Related resources
Object storage system IBM Cloud Object Storage	Access pools Accesser nodes Mirrors Sites Slicestor nodes Storage pools Vaults	Accounts Containers	Back-end storage systems Fabrics GPFS clusters Hypervisors Servers Switches

Switches

A *switch* is a network infrastructure component to which multiple nodes attach. In IBM Spectrum Control, you can monitor Fibre Channel (FC) switches. The following table shows the internal resources and related resources that you can monitor for switches.

Table 33. Switch resources that you can monitor			
Switch icon	Internal resources	Related resources	*Connected resources (to ports)
Physical switch Logical switch	✓ Inter-switch connections ■ Blades ■ Logical switches ■ Physical switches ■ Ports	Fabrics Hypervisors Servers Storage systems Switches Discovered ports	All monitored resources that are connected to ports, plus the following resources: Unmanaged resource. The name is determined from a fabric zone alias. Unmanaged resource. The name is determined from a storage system's mapping to a host. Unmanaged resource. The name is determined from one of a storage system's mappings to a host when multiple mappings are found.

*Connected resources (to ports): The icons for any resources that are managed by IBM Spectrum Control and connected to a port are shown in this column. Icons are also shown for connected resources whose names are known from a fabric zone alias or a storage system's mapping to a host (host connection) but are not currently being managed. If both a fabric zone alias and a storage system's mapping to a host are found for the same unmanaged resource, the fabric zone alias is shown as the name.

Related tasks

Adding ESX and ESXi hypervisors

Add ESX and ESXi hypervisors for monitoring by IBM Spectrum Control.

Adding storage systems

Add the storage systems that you want to monitor in IBM Spectrum Control.

Adding fabrics and switches

Add fabrics and switches for monitoring.

Adding resources

You must add resources to IBM Spectrum Control before you can collect data, generate reports, and manage storage that is related to those resources.

Related concepts

Alerting

Specify conditions that trigger alerts and the actions to take when those alerts are triggered, such as notify an email address. Use alert policies to define those alert conditions and notification settings for a group of resources.

Related reference

Triggering conditions for alerts

Define alerts so that IBM Spectrum Control automatically notifies you when certain conditions or events are detected on monitored resources. Such conditions are the *triggering conditions* for the alert. The specific conditions that can trigger alerts depend on the type of resource that is being monitored.

Required user roles for monitoring resources

When you add storage systems, switches, and hypervisors for monitoring, you must provide a user name and password for logging in to those resources. The role or user group that is assigned to the user name determines the data collection and storage functions that you can use.

The following roles are associated with the user names that IBM Spectrum Control uses to log in to resources. Specify user names when you add a resource for monitoring. These roles are different from IBM Spectrum Control roles, which are assigned to users that log in to IBM Spectrum Control.

Any roles that are not listed, but include the privileges of the roles that are listed, can also be used for monitoring resources.

Table 34. Required roles for storage system, switch, and hypervisor users			
Resource	Required role for performance monitors	Required role for probes	Required role for provisioning, optimizing, and transforming
DS6000 DS8000	Monitor role or higher	Monitor role or higher	Administrator, Physical operator, or Logical operator
SAN Volume Controller	For versions earlier than 8.3.1.2, Administrator, or SecurityAdmin. For version 8.3.1.2 or later, any role but some limitations might apply.	Monitor role or higher	Administrator
Storwize V3500 Storwize V3700 Storwize V5000 Storwize V7000	For versions earlier than 8.3.1.2, Administrator, or SecurityAdmin. For version 8.3.1.2 or later, any role but some limitations might apply.	Monitor role or higher	Administrator
Storwize V7000 Unified (block storage)	Administrator	Monitor role or higher	Administrator
FlashSystem 5000 FlashSystem 5100 FlashSystem 7200 FlashSystem 9100 FlashSystem 9200 FlashSystem V840 FlashSystem V9000	For versions earlier than 8.3.1.2, Administrator, or SecurityAdmin. For version 8.3.1.2 or later, any role but some limitations might apply.	Monitor role or higher	Administrator

Table 34. Required roles for storage system, switch, and hypervisor users (continued)			
Resource	Required role for performance monitors	Required role for probes	Required role for provisioning, optimizing, and transforming
IBM Spectrum Virtualize for Public Cloud	For versions earlier than 8.3.1.2, Administrator, or SecurityAdmin. For version 8.3.1.2 or later, any role but some limitations might apply.	Monitor role or higher	Administrator
Storwize V7000 Unified (file storage)	Performance monitoring is not available	Monitor	Administrator, or Export Administrator + Storage Administrator.
XIV	Monitor role or higher	Any role	Storage administrator
IBM Spectrum Accelerate IBM FlashSystem A9000 IBM FlashSystem A9000R	Monitor role or higher	Any role	Not available
IBM FlashSystem 840 IBM FlashSystem 900	Monitor role or higher	Monitor role or higher	Not available
IBM Spectrum Scale (file storage)	Any role	Root or non-root*	Not available
IBM Spectrum Scale (object storage)	Performance monitoring is not available	The Keystone admin role. Information is collected only for the object storage accounts and containers that the user has access to. If you want to monitor all accounts and containers, the user must also be assigned the role that is defined in the reseller_admin_role configuration option in the Swift proxy server. The default value for the reseller_admin_role option is ResellerAdmin.	Not available

Table 34. Required roles for storage system, switch, and hypervisor users (continued)			
Resource	Required role for performance monitors	Required role for probes	Required role for provisioning, optimizing, and transforming
IBM Cloud Object Storage	Performance monitoring is not available	Operator, System Administrator, or Super User role	Not available
Dell EMC storage systems	Operator role or higher	Operator role or higher	Not available
Hitachi storage systems	In Hitachi Device Manager: The user for the Export Tool must have the Storage Administrator (Performance Management) role	In Hitachi Command Suite: View Group or higher	Provisioning is not available
NetApp storage systems	Operator role or higher	Operator role or higher	Not available
Pure storage systems	Role with read-only permission or higher	Role with read-only permission or higher	Not available
Cisco switches	Network-admin	Network-admin	Not applicable
Brocade switches with Fabric OS 8.2.1 or later	User or admin role that has the chassis-role permission	User or admin role that has the chassis-role permission	Not applicable
Brocade switches with a Fabric OS version earlier than 8.2.1	Administrator user in Brocade Network Advisor (BNA)	Administrator user in BNA	Not applicable
Hypervisors (such as ESX, ESXi, and vCenter Servers)	Performance monitoring is not available	A role that has permission to browse through data stores. For example: Administrator role or Virtual Machine Power® User role.	Any role
Rollup servers	Performance monitoring is not available	Administrator	Not available

Note: * You can add IBM Spectrum Scale and GSS systems as a non-root user, but that user must have privileges to run a set of specified administration commands using the sudo command on the cluster node.

Adding storage systems

Add the storage systems that you want to monitor in IBM Spectrum Control.

When block storage systems are added, asset, capacity, and configuration metadata and performance metadata is automatically collected and analyzed. When storage systems that manage block and file

storage are added, the asset, capacity, and configuration metadata for file storage is also automatically collected and analyzed.

Metadata Collection Schedule: By default, asset, capacity, and configuration metadata is aggregated and collected daily. Performance metadata is collected at the following intervals:

- Every 5 minutes for IBM block storage systems
- Every 5 minutes for Dell EMC Unity storage systems
- Every 5 minutes for Hitachi VSP storage systems
- Every 5 minutes for NetApp storage systems that are running ONTAP 9
- Every 5 minutes for Pure storage systems
- Every 15 minutes for other Dell EMC block storage systems.

You can add multiple storage systems at the same time, or add them one by one. To add multiple storage systems, the storage systems must be of the same type and share authentication credentials.

To add storage systems, you must have Administrator privileges.

You can add the following types of storage systems for monitoring:

Table 35. Storage systems that can be monitored in IBM Spectrum Control

Storage system	Can add many at same time.
DS8000	✓
DS4000, DS5000, or DS6000	
Dell EMC Unity	✓
Dell EMC VMAX	
Dell EMC VNX	
Dell EMC VNXe	
FlashSystem 5000	✓
FlashSystem 5100	✓
FlashSystem 7200	✓
FlashSystem 9100	✓
FlashSystem 9200	✓
FlashSystem V840	✓
FlashSystem V9000	✓
FlashSystem 840	~
FlashSystem 900	~
FlashSystem A9000	✓
FlashSystem A9000R	✓
Hitachi VSP	✓
NetApp ONTAP 9	✓
NetApp ONTAP 8 (7- mode)	
Pure	~
IBM Spectrum Accelerate	~
IBM Spectrum Scale (ESS and GSS)	
IBM Spectrum Virtualize	✓

Table 35. Storage systems that can be monitored in IBM Spectrum Control (continued)

Storage system	Can add many at same time.
IBM Spectrum Virtualize for Public Cloud	✓
SAN Volume Controller	✓
Storwize V3500	✓
Storwize V3700	✓
Storwize V5000	✓
Storwize V5000E	✓
Storwize V7000	✓
IBM Storwize Flex System V7000 Storage Node	
Storwize V7000 Unified	
XIV	✓
All others (managed by SMI-S providers)	

The following storage systems run IBM Spectrum Virtualize to virtualize their storage: SAN Volume Controller, IBM Spectrum Virtualize for Public Cloud, IBM Spectrum Virtualize as Software Only, IBM Storwize family, FlashSystem 5000, FlashSystem 5100, FlashSystem 7200, FlashSystem 9100, FlashSystem 9200, FlashSystem V840, and FlashSystem V9000. In this documentation, IBM Spectrum Virtualize is used to refer collectively to IBM SAN Volume Controller, IBM Spectrum Virtualize for Public Cloud, IBM Spectrum Virtualize as Software Only, and IBM Storwize storage systems, and to IBM FlashSystem devices that run IBM Spectrum Virtualize.

Tips:

- To monitor the performance of a Storwize V7000 Unified storage system, you must add it as a block storage system.
- To monitor an IBM Spectrum Virtualize for Public Cloud storage system, you must configure it for communication with IBM Spectrum Control. For more information, see Configuring IBM Spectrum Virtualize for Public Cloud.
- To add block storage systems, go to Storage > Block Storage Systems. To add a file storage system, go to Storage > File Storage Systems. To add an object storage system, go to Storage > Object Storage Systems.
- 2. Click Add Storage Systems or Add Storage System.
- 3. Click the icon for the type of storage system that you want to add.

Click this icon	To add these storage systems
Storwize Family	IBM® Storwize® V3500, Storwize V3700, Storwize V5000, Storwize V5000E, Storwize V7000
SVC or Spectrum Virtualize	SAN Volume Controller, IBM Spectrum Virtualize for Public Cloud, or IBM Spectrum Virtualize software-only clusters
FlashSystem Family	FlashSystem 5000 (formerly known as IBM Storwize V5000), FlashSystem 5100, FlashSystem 7200 (formerly known as IBM Storwize V7000), FlashSystem 9100, FlashSystem 9200, FlashSystem V840, FlashSystem V9000, FlashSystem 840, FlashSystem 900, FlashSystem A9000R
Dell EMC	UnityVMAX family with support for SMI-S 1.6

Click this icon	To add these storage systems	
	 VNX family with support for SMI-S 1.6 VNXe family with support for SMI-S 1.6 All other Dell EMC storage systems with support for SMI-S 1.2 	
Hitachi	Hitachi VSP	
NetApp	• ONTAP 9 • ONTAP 8 (7- mode)	
Pure	FlashArray//M FlashArray//X	
Others	Other non-IBM storage systems	

4. Specify the connection details and the authentication credentials for the storage systems that you want to add, then click **Connect**.

You can close the window and continue your work. To change the default probes and performance monitors, keep the window open and update them when the storage systems are connected.

The storage systems are added for monitoring. Each storage system is automatically added to the default alert policy for the storage system type. A probe is automatically run to collect status and asset information about the storage systems.

Storage systems can be configured for block storage, file storage, a combination of block and file storage, or a combination of file and object storage. The method that a storage system uses for managing data determines the internal resources and object resources that are monitored.

After you add storage systems, a task is added to the list of running tasks on the page banner. If the task fails, the task is added to the list of failed tasks on the page banner. You can retry or cancel tasks that failed.

When the collection of data is complete, you can view status information and capacity data about the storage system on the resource list page for the block, file, or object storage system. For example, for block storage systems, you can view the information on the **Block Storage Systems** page.

After the probe is run, agentless servers are updated with any new information about the host connections on the storage systems. You can review the agentless servers on the **Servers** page, and make any changes that are required.

Managing alert conditions and notification settings: Review the alert configuration for the storage system. You can change which alert policy manages the storage system. You can also set a storage system to not be managed by any policy.

Hitachi

Add Hitachi Virtual Storage Platform (VSP) F and G Series storage systems to get performance, asset, capacity, and configuration metadata analyzed for block storage so that you can detect performance issues, changes in storage usage, and plan for future storage needs.

You can add a VSP system by connecting to the Hitachi Command Suite that is managing the device. By default, the asset, capacity, and configuration metadata that is collected from storage systems is refreshed every 24 hours. Performance metadata is collected at a default collection interval of 5 minutes.

Use the following information to add VSP storage systems so that they can be monitored and their metadata can be collected, analyzed, and presented in the GUI. Add a VSP storage system by connecting to the Hitachi Command Suite that manages the device. Connect to the Hitachi Device Manager to enable the Export Tool to collect performance information from the storage system.

Supported versions: To view the versions of Hitachi VSP storage systems that are supported in IBM Spectrum Control, go to the Hitachi support page.

Restriction: Performance monitoring is not supported when IBM Spectrum Control is installed on AIX.

Complete the following steps to add your Hitachi VSP storage systems for monitoring.

- Add Hitachi VSP F and G Series by connecting to Hitachi Command Suite. See <u>"Credentials for Hitachi Command Suite"</u> on page 122.
- Enable performance monitoring while you are adding your Hitachi VSP systems. See <u>"Performance" on page 122.</u>
- Install the Hitachi Export Tool to collect performance information and export it from your storage system. See "Installing the Hitachi Export Tool" on page 123.

Add Hitachi VSP F and G Series

When you select Hitachi VSP, you can add one or more storage systems by entering the connection information for the Hitachi Command Suite that is managing those devices.

Credentials for Hitachi Command Suite

IBM Spectrum Control monitors capacity and configuration information about Hitachi devices by using the Hitachi Command Suite. Provide the connection information for the Hitachi Command Suite installation that is monitoring the devices.

Host name or IP address

The host name or IP address that you use to connect to Hitachi Command Suite. Depending on what is supported in your environment, you can enter an Internet Protocol version 4 (IPv4) or IPv6 address. If you enter an IPv6 address, the preferred representation is written as eight groups of four hexadecimal digits. Example: 2001:DB95:0000:1234:0000:0000:5678:ABCD.

User Name and Password

The account credentials that are required to connect to Hitachi Command Suite. The role or user group that is assigned to the user name must have the appropriate privileges to monitor the data that is collected and, if required, to change the frequency of the data collection schedules. <u>Learn more</u> about the role requirements for the user name.

Performance

Enable performance monitoring

When you enable performance monitoring in IBM Spectrum Control, the user name and password that you provided to connect to the Hitachi Command Suite are also used to connect to Hitachi Device Manager. Hitachi Device Manager provides the connection information for the storage system to the Export Tool, which then sends the performance information to IBM Spectrum Control

Requirement: You might already have created credentials for Hitachi Device Manager, but you must create new credentials to match the credentials used for Command Suite. Device Manager uses these credentials exclusively for the Export Tool.

Important: To collect performance metadata, Hitachi requires that the password for connecting to the storage system be temporarily stored as clear text in a file on the server where IBM Spectrum Control is installed. For security reasons, ensure access to that server is restricted to key personnel. The file is automatically deleted after metadata is collected.

User Name and Password

The account credentials that are required to connect to Hitachi Device Manager. The role or user group that is assigned to the user name must have the appropriate privileges to monitor the data that is collected and, if required, to change the frequency of the data collection schedules. <u>Learn more</u> about the role requirements for the user name.

Capacity

Asset, capacity, and configuration metadata is collected by probes. By default, a probe collects metadata from Hitachi VSP systems once every 24 hours.

Configure

If you want to change the default probes and performance monitors, keep the **Add Storage System** window open. When the storage system or systems are connected, select the systems that you want to configure. To select multiple storage systems, press Ctrl and click or Shift and click. The default configuration settings are applied to the remaining storage systems.

Display Name

The name of the storage system that is displayed in the IBM Spectrum Control interface. If you do not enter a value, the default name is provided. This is the name that is provided with the storage system.

Location

The location, such as the geographical location or the building, where the resource is located.

Data Collection

Schedule the probe and performance monitor for the storage system. The probe collects status, asset, and storage information about the storage system. The performance monitor collects metrics that measure the performance of the storage system.

Probe

Enter the time and schedule for the storage system probe to run.

Enter the time in *hh:mm* format, where *hh* equals the hour and *mm* equals the minute. The time zone that is shown is determined by the location of the Data server for IBM Spectrum Control.

Performance Monitor

If you enabled performance monitoring when you added the Hitachi storage systems, you can change the performance monitoring interval from every 5 minutes to every 60 minutes. You can also disable performance monitoring.

Installing the Hitachi Export Tool

To collect performance data from a Hitachi storage system, you must install the Hitachi Export Tool on the server or virtual machine where IBM Spectrum Control is installed. If you have many Hitachi storage systems in your environment, you must install the appropriate version of the Export Tool for each model of storage system.

To collect performance metadata, the Hitachi Export Tool requires that the password for connecting to the storage system is temporarily stored as clear text in a file on the server or virtual machine where IBM Spectrum Control is installed. For security reasons, ensure access to that server or virtual machine is restricted to key personnel. The file is automatically deleted after metadata is collected.

1. Use your Hitachi account to download the Export Tool from the following location: https://knowledge.hitachivantara.com/Knowledge/Storage/
How_to_Download_the_Appropriate_Export_Tool_Version_Specific_to_Array_Microcode

Tip: If you have more than one model of Hitachi storage system, you must download the version of the Export Tool for each model.

2. To install the tool into the default directory, first create the following directory structure where hitachi_model is the directory name that corresponds to the Export Tool for the model of your storage system.

Windows

Program Files\monitor\export\hitachi_model

Linux

/opt/monitor/export/hitachi_model

Where hitachi_model is the directory name. See the following table for the directory details.

Table 36. Hitachi storage system models and export tool directories		
Model	Property Name	Directory Name
Hitachi Virtual Storage Platform 5100/5100H/5500/5500H	hitachi_5000_ExportToolP ath	5000
Hitachi Virtual Storage Platform E990	hitachi_E_ExportToolPath	eseries
Hitachi VSP G200/400/600/800 and F400/600/800	hitachi_FGX00_ExportTool Path	fgx00
Hitachi VSP G130/350/370/700/900 and F350/370/700/900	hitachi_FGXX0_ExportTool Path	fgxx0
Hitachi VSP G1000/1500 and F1500	hitachi_FG1X00_ExportToo lPath	fg1x00
Hitachi VSP N400/600/800	hitachi_N_ExportToolPath	nseries
Hitachi Virtual Storage Platform (VSP)	hitachi_LegacyVSP_Export ToolPath	legacyvsp

- 3. Ensure that the Export Tool directory has write permission for users.
- 4. For each Export Tool that you downloaded, open the disk image and extract the tool for your chosen operating system.
- 5. Install the tool into the default directory.
- 6. Optional: If it is not possible to run the Export Tool from the default directory, another directory can be used. However, you must modify the setup.properties file in the DataCollector/conf directory to specify the new location.

Modify the hitachi_model_ExportToolPath property for your specific Hitachi storage system as in the following examples for Hitachi VSP G200/400/600/800 and F400/600/800.

Windows)

Linux

hitachi_FGX00_ExportToolPath=/usr/abc/HitachiExportTool/export/fgx00

7. Optional: If you changed the setup.properties file, save it and then restart the device server.

Adding and configuring NetApp resources

After adding and configuring a NetApp resource, you can use IBM Spectrum Control to administer and monitor the status, configuration, capacity, performance, and information that is collected about the resource.

A NetApp ONTAP 9 resource can be added to IBM Spectrum Control as either a block or file storage system to view all of the block and file data on the system. A NetApp ONTAP 8 (7- mode) resource can be added to IBM Spectrum Control as a block storage device, a file storage device, or both. IBM Spectrum Control collects different information from the device, depending on the type of storage system you specify. For all NetApp storage systems IBM Spectrum Control collects data by directly connecting to the resource.

Important:

• For new IBM Spectrum Control installations, you can no longer deploy Storage Resource agents for monitoring NetApp resources. You must now monitor those resources through an SMI-S provider or a direct connection. However, if you upgrade from a previous version of the product, you can continue to manage NetApp resources that have Storage Resource agents deployed from that previous version.

- When you upgrade from a previous version to IBM Spectrum Control 5.2.10 or later, the existing alert definitions for NetApp filers are migrated to the new version. However, to generate alerts for those filers, you must monitor them through direct connections rather than through Storage Resource agents. To set up direct connections for those NetApp filers, add them again for monitoring and enter the credentials for connecting to them directly. After you add the NetApp filers, their previous entries are automatically removed and alerts will be generated for the alert definitions that were migrated.
- For a complete list of NetApp devices that can be monitored, see NetApp support page.
- For NetApp storage systems managed by SMI-S providers, performance monitoring is supported only for volumes and front-end ports. It is not supported at the storage system level.

To add a NetApp resource for monitoring by IBM Spectrum Control, complete the following steps:

- 1. To add a NetApp resource as a block storage system, go to **Storage** > **Block Storage Systems**. To add the resource as a file storage system, go to **Storage** > **File Storage Systems**.
- 2. Click Add Storage System.
- 3. Click the **NetApp** icon to add the resource as a storage system.
- 4. Complete the connection information for the storage system and schedule the collection of data.

After the initial discovery has completed, you can view information about the NetApp device in the **File Storage Systems** or **Block Storage Systems** pages in the GUI.

Antivirus software restriction: If your antivirus software is set on the maximum mode, it might prevent your ONTAP 9 storage systems from being added to IBM Spectrum Control. For more information about how to configure your antivirus software, see Installation checklists for IBM Spectrum Control.

By default, asset, capacity, and configuration metadata is aggregated and collected daily. You can schedule daily capacity and inventory reports to gain insights about your NetApp storage systems.

Performance metadata is collected for NetApp storage systems running ONTAP 9 at collection intervals of 5 minutes and 60 minutes. The default interval is 5 minutes. For other NetApp storage systems, performance metadata is collected every 15 minutes.

Configuring IBM Spectrum Virtualize for Public Cloud for monitoring

Configure IBM Spectrum Virtualize for Public Cloud for monitoring with IBM Spectrum Control.

IBM Spectrum Virtualize is a software-defined storage solution that has been proven for years in SAN Volume Controller and the IBM Storwize family. IBM Spectrum Virtualize for Public Cloud extends that solution to a hybrid-cloud or cloud-based model, where servers, storage, and network infrastructure are delivered in a public cloud environment. It can be deployed on either IBM® Cloud or Amazon Web Services (AWS) cloud infrastructures.

With IBM Spectrum Control, you can view the capacity, space usage, and performance of your IBM Spectrum Virtualize for Public Cloud storage systems. Other monitoring features, such as alerting, health checking, advanced analytics, and reporting are also supported.

Before you can add an IBM Spectrum Virtualize for Public Cloud storage system for monitoring, you must ensure that IBM Spectrum Control can connect to it. To enable a connection, you can use the site-to-site VPN IPSec tunnel that exists between the on-premises environment and the IBM Spectrum Virtualize for Public Cloud storage systems.

Monitoring IBM Spectrum Virtualize for Public Cloud with on-premises data collection (Site to Site VPN IPsec)

You can connect and monitor IBM Spectrum Virtualize for Public Cloud storage with IBM Spectrum Control by using the site-to-site virtual private network (VPN) IPSec tunnel that exists between the onpremises environment and the IBM Spectrum Virtualize for Public Cloud instances.

The VPN IPSec site-to-site tunnel is a secure communication network between the cloud infrastructure and the on-premises environment. Network communication between the private subnets is controlled by the access control list (ACL) that is populated when you create the VPN IPSec site-to-site tunnel.

Typically, a bi-directional, IPsec site-to-site tunnel is limited to the subnets that contain the following IP addresses:

- On-premises IBM Spectrum Virtualize cluster and replication target
- Cloud-based IBM Spectrum Virtualize for Public Cloud cluster and replication target

To use the IPsec site-to-site tunnel for communication between IBM Spectrum Control and IBM Spectrum Virtualize for Public Cloud, you must include the IP addresses of your IBM Spectrum Control servers (or the subnet in which they reside) in the tunnel definition as one of the on-premises endpoints.

For example, in the AWS Management Console, you can define the following:

- The external (internet-routable) IP address of the on-premises IPsec tunnel endpoint.
- The subnet / IP addresses for the on-premises IBM Spectrum Controlservers, IBM Spectrum Virtualize cluster, and replication target that will communicate through the tunnel.
- The IP addresses of the IBM Spectrum Virtualize for Public Cloud cluster and replication target in AWS that will communicate through the tunnel.

When defined, you can export a configuration file that can be used in different IPsec VPN devices, such as Vyatta and Juniper (Junos VPN Site Secure).

Contact your network and firewall administrators to help set up this communication between IBM Spectrum Control and IBM Spectrum Virtualize for Public Cloud.

After configuring the connection between IBM Spectrum Virtualize for Public Cloud and IBM Spectrum Control, add the storage system for monitoring by completing the following steps:

- 1. In IBM Spectrum Control, go to Storage > Block Storage Systems.
- 2. Click Add Storage Systems.
- 3. Click the **SVC or Spectrum Virtualize** icon.
- 4. Specify the IP address and authentication credentials for the IBM Spectrum Virtualize for Public Cloud instance that you want to monitor.
- 5. Click Connect.

The storage system is added for monitoring and is automatically added to the default alert policy for the storage system type. Data collection is automatically run to collect status, configuration, capacity, and performance metadata about the storage system.

By default, asset, capacity, and configuration metadata is aggregated and collected daily. Performance metadata is collected every 5 minutes. You can schedule daily capacity and inventory reports to gain insights about your IBM Spectrum Virtualize for Public Cloud storage systems.

User roles for collecting performance metadata from IBM Spectrum Virtualize

When you add IBM Spectrum Virtualize storage systems, you must add the credentials of a user who can log on to the storage system and collect performance metadata.

The required role of the user on your storage system for collecting performance metadata depends on the version of IBM Spectrum Virtualize.

Versions earlier than 8.3.1.2

The user must have the role of Administrator or SecurityAdmin.

Version 8.3.1.2 or later

The user can have any role, such as the Monitor role.

Important information about user roles for IBM Spectrum Virtualize version 8.3.1.2 or later

If the collection of performance metadata is stopped on the storage system, the role of the storage system user determines whether the collection can be automatically restarted. For example, if the storage system user has the Monitor role and your administrator has stopped the collection on the storage system, the collection must be manually restarted by your administrator before IBM Spectrum Control can collect performance metadata.

User has privileges to start collections

If the user has the Administrator or SecurityAdmin role on the storage system, the collection is automatically restarted and the performance metadata is collected.

User does not have privileges to start collections

If the user has the Monitor role on the storage system, IBM Spectrum Control can't automatically restart the collection. To manually start the collection of performance metadata on the storage system, complete the following steps:

- 1. Log on to the storage system as a user with the Administrator, or SecurityAdmin role.
- 2. Run the following command:

```
svctask startstats -interval time in minutes
```

The interval must be less than or equal to the performance monitor interval for the storage system in IBM Spectrum Control. To check the value, in the IBM Spectrum Control GUI, go to the **Block Storage Systems** page and view the value for Performance Monitor Interval for the storage system.

Monitoring IBM Spectrum Scale without requiring root privileges

Enable users that do not have root privileges on a IBM Spectrum Scale cluster node to monitor IBM Spectrum Scale storage systems.

Before you add the IBM Spectrum Scale storage system for monitoring, complete this task.

To grant access to the user, make the following changes to the sudoers file on the cluster node that is used for authentication:

- Add the set of administration commands that are required to monitor the storage system.
- Associate the user with the set of administration commands.
- 1. Log on to the cluster node that is used for authentication with a user name that has root privileges.
- 2. To edit the sudoers file, enter the following command:

```
visudo -f /etc/sudoers
```

3. Add the following command aliases to the sudoers file.

The command aliases contain the commands that the user must be able to issue to monitor the storage system.

Important: Each command alias must be on a single line without line breaks.

```
Cmnd_Alias TPC_GPFS_MMCMD = /usr/lpp/mmfs/bin/mmsdrquery, /usr/lpp/mmfs/bin/mmlsconfig,
    /usr/lpp/mmfs/bin/mmgetstate, /usr/lpp/mmfs/bin/mmlsnodeclass, /usr/lpp/mmfs/bin/mmlsfs,
    /usr/lpp/mmfs/bin/mmdf, /usr/lpp/mmfs/bin/mmlsnsd, /usr/lpp/mmfs/bin/mmlsfileset,
      /usr/lpp/mmfs/bin/mmcloudgateway, /usr/lpp/mmfs/bin/mmlsmount, /usr/lpp/mmfs/bin/
mmlssnapshot
      /usr/lpp/mmfs/bin/mmrepquota, /usr/lpp/mmfs/bin/mmlspolicy, /usr/lpp/mmfs/bin/
mmapplypolicy
Cmnd_Alias TPC_GPFS_MMDSH = /usr/lpp/mmfs/bin/mmdsh -N * /usr/lpp/mmfs/bin/mmdiag --version,
     /usr/lpp/mmfs/bin/mmdsh -N * /lib/udev/scsi_id --whitelisted *, /usr/lpp/mmfs/bin/mmdsh -N * /sbin/blockdev --getsize64 *, /usr/lpp/mmfs/bin/mmdsh -N * /usr/bin/getconf DISK_SIZE *,
      /usr/lpp/mmfs/bin/mmdsh -f 20000 -N linuxNodes 'cat /sys/class/fc_host/*',
      /usr/lpp/mmfs/bin/mmdsh -N * /usr/lpp/mmfs/bin/mmces node list,
/usr/lpp/mmfs/bin/mmdsh -N * /usr/lpp/mmfs/bin/mmces service list -a,
      /usr/lpp/mmfs/bin/mmdsh -N * /usr/lpp/mmfs/bin/mmces address list|grep
object_database_node,
      /usr/lpp/mmfs/bin/mmdsh -N * /usr/lpp/mmfs/bin/mmces address list
       --by-node|grep object_database_node,
     /usr/lpp/mmfs/bin/mmdsh - - N cesNodes /usr/lpp/mmfs/bin/mmobj config list
--ccrfile object-server.conf --section DEFAULT --property devices,
/usr/lpp/mmfs/bin/mmdsh -f 200000 -v -N * "test -e /opt/IBM/zimon/ZIMonSensors.cfg &&
      (grep -w collectors -A 4 /opt/IBM/zimon/ZIMonSensors.cfg | grep -w host) || true<sup>m</sup>
      /usr/lpp/mmfs/bin/mmdsh -f 20000 -v -N nonWindowsNodes hostname
Cmnd_Alias TPC_GPFS_OTHER = /bin/cat *release, /usr/bin/lsb_release -a
Cmnd_Alias TPC_GPFS_CMDS = TPC_GPFS_MMCMD, TPC_GPFS_MMDSH, TPC_GPFS_OTHER
```

4. To enable the user to issue the commands, add the following lines after the command aliases:

```
Defaults: user\_name !requiretty
```

```
user_name ALL=(ALL) TPC_GPFS_CMDS
```

where user_name is the user name that you enter when you add the storage system for monitoring.

The user that you added to the sudoers file can now monitor the IBM Spectrum Scale storage system.

Verifying that asset, capacity, and configuration metadata can be collected for object storage
To probe object storage, IBM Spectrum Control server must be able to connect to the OpenStack Swift
and Keystone endpoints that are used to access object services.

To verify that the server can access the object services, complete these steps:

1. List the URLs for the Keystone and Swift services by running commands on a IBM Spectrum Scale cluster node that is configured for object storage. Run the following commands with a user name that has root privileges:

```
. ~/openrc openstack endpoint list
```

2. Ensure that the IBM Spectrum Control server can connect to the IP addresses and host names that are included in the Keystone and Swift services URLs.

For example, the URL for the Keystone service might be http://gpfs420proto1:5000/v3. The IBM Spectrum Control server must be able to connect to the gpfs420proto1 host name.

Configuring the collection of performance data for IBM Spectrum Scale

Configure the collection of information about the performance of IBM Spectrum Scale file storage systems.

To collect performance data, you must first configure the IBM Spectrum Scale performance monitoring tool on the IBM Spectrum Scale cluster. Then, use the IBM Spectrum Control GUI to schedule the collection of performance data.

1. Configure the sensor component of the IBM Spectrum Scale performance monitoring tool by editing the /opt/IBM/zimon/ZIMonSensors.cfg file on the sensor nodes.

Set the host property in the collectors section of the file to the IBM Spectrum Scale cluster node where the collector component will run. The host property must be set to one of the following options:

- An IP address that can be reached by the IBM Spectrum Control server
- A host name that resolves to an IP address that is reachable by the IBM Spectrum Control server
- 2. Enable the following sensors on each node for which you want to collect performance data: CPU, Memory, GPFSNode, GPFSFileSystem.
- 3. Configure the collector component of the performance monitoring tool so that the IBM Spectrum Control server can connect to port 9084 on the cluster node where the collector component will run.
 - Configure the collector component by editing the /opt/IBM/zimon/ZIMonCollector.cfg file on the IBM Spectrum Scale cluster node where the collector component will run. Ensure that the queryinterface property is set to "0.0.0.0".
- 4. Start the collector component.
- 5. Start the sensor component of the performance monitoring tool on each node for which you want to collect performance data.
 - By default in IBM Spectrum Scale version 5.0.0 or higher, the sensors are started on all nodes.
- 6. Schedule the collection of performance data by creating a performance monitor in the IBM Spectrum Control GUI.

You can create performance monitors in the GUI when you add resources for monitoring or you can create them later.

A performance monitor is created for the IBM Spectrum Scale storage system. If a successful probe run is completed for the resource, the performance monitor runs according to the defined interval.

To check the progress of a performance monitor, you can complete the following actions:

- View the Performance Monitor Status column on the **File Storage Systems** page.
- From the File Storage Systems page, right-click a row, and select Data Collection > Open Performance Monitor Logs.

To learn more about configuring the performance monitoring tool, see https://www.ibm.com/support/knowledgecenter/STXKQY_5.0.1/com.ibm.spectrum.scale.v5r01.doc/bl1adv_PMToverview.htm.

Configuring OpenStack access to monitor the object storage system

Configure OpenStack access for the user name that is used to monitor the IBM Spectrum Scale object storage system.

Before you add the IBM Spectrum Scale object storage system for monitoring, ensure that the user name that you use to monitor the storage system has access to object storage accounts.

Tip: The terms "account" and "project" mean the same thing. Swift, the OpenStack object storage service, uses the term "account" and Keystone, the OpenStack identity service, uses the term "project."

- 1. Set the object storage account and the domain for the user. By default, the domain is set to Default and cannot be modified after you create the user account. Choose one of the following actions:
 - a. To set the account and domain when you create a user account, use the following command:

```
openstack user create --domain domainname --project projectname --password Password username
```

b. To set the account and the domain for an existing user name, use the following command:

```
openstack user set --project projectname username
```

2. Assign the admin role for an object storage account to the user name. Use the following command:

```
openstack role add --user username --project projectname admin
```

3. To monitor all accounts on the object storage system, assign to the user name the role that is defined in the **reseller_admin_role** configuration option in the Swift proxy server. The default value for the **reseller_admin_role** option is ResellerAdmin.

For example, use the following command:

```
openstack role add --user username --project projectname ResellerAdmin
```

Restriction: If you do not assign the ResellerAdmin role, information is collected only for the object storage accounts that the user has admin access to.

To learn more about configuring and assigning user roles in OpenStack, see https://docs.openstack.org/.

Adding fabrics and switches

Add fabrics and switches for monitoring.

When you add switches for monitoring, you must provide a user name and password for logging in to those switches. The role or user group of the user determines the type of data that can be collected. To see the privileges that are required for switch users, see <u>"Required user roles for monitoring resources" on page 116.</u>

For a Cisco switch to successfully receive and respond to queries from IBM Spectrum Control, the following basic requirements must be met:

• IBM Spectrum Control can use SNMPv3 (preferred) or SNMPv1 to probe switches and fabrics. The SNMPv3 protocol is preferred because it provides better security, but switches that use the SNMPv1 protocol are also supported. Some switches are configured to use SNMPv3 by default.

- The Fibre Alliance FC Management MIB (FA MIB) and Fibre Channel Fabric Element MIB (FE MIB) must be enabled on the switch.
- When you use the SNMPv1 protocol, the community string that is configured in IBM Spectrum Control must match one of the community strings that are configured on the switch with read access. Additionally, Cisco switches must have a community string match for write access. The default community strings in IBM Spectrum Control are "public" for read access and "private" for write access. Other community strings can be defined on the switches, but are not used. Community strings are not relevant when you use the SNMPv3 protocol.
- SNMP access control lists must include the IBM Spectrum Control system. These access control lists are defined and configured on the switches. Some lists automatically include all hosts, while others exclude all by default.
- The Fibre Channel (FC) or Fibre Channel over Ethernet (FCoE) protocols must be enabled on the switch.
 Some switches, such as the Cisco Nexus 5000 series, require you to enable these protocols. Otherwise,
 IBM Spectrum Control does not recognize the switch when you try to add it using the Add Switches and
 Fabrics for Monitoring dialog. For instructions on how to configure Cisco switches for FCoE
 enablement, go to the Cisco product website at http://www.cisco.com and click Support.

IBM Spectrum Control can monitor Brocade switches and fabrics in either of the following ways:

- Use the REST API to manage the switches and fabrics.
- Use the Brocade Network Advisor (BNA) Storage Management Initiative (SMI) agent to manage the switches and fabrics. The SMI agent is embedded in the BNA. BNA can manage multiple fabrics within and across data centers. When you configure BNA, you set up one switch to be the *seed* switch that interconnects to all the other switches in the fabric.

Important: The embedded SMI agent is only available in the Professional Plus and Enterprise editions of BNA. BNA is no longer available for purchase, although it will be supported by Broadcom until February 2022. For more information about end of support for BNA, see https://www.broadcom.com/support/fibre-channel-networking/eol.

In IBM Spectrum Control, you can add switches and fabrics for monitoring at the same time. You can add the following types of switches:

- Brocade
- Cisco

For a complete list of the firmware levels of supported switches, see https://www.ibm.com/support/pages/ibm-spectrum-control-53x-switches.

- 1. From the Network menu, click Switches or Fabrics.
- 2. Click Add Switch or Add Fabric.
- 3. Select the vendor: Brocade or Cisco.
- 4. Specify how you want to discover the switches and the fabrics that are connected to the switches. For Brocade switches and fabrics, you can connect using either REST API or an SMI-S provider on Brocade Network Advisor (BNA). For Cisco switches and fabrics, you must use SNMP agents.

Type of switch	Steps
Brocade	Complete the following steps:
	a. Select the version of the Fabric OS that is running on the switches.
	b. If you select 8.2.1 or later , specify the connection details and the authentication credentials for the switches that you want to add. You need to add only one switch from each fabric to IBM Spectrum Control. If other switches that are on the same fabric are running Fabric OS 8.2.1 or later and have the same username, password, protocol, and port, they are added to IBM Spectrum Control automatically.

Type of switch	Steps
	The user must have a user or admin account with chassis-role privileges for the switches that are added.
	Click Next.
	c. If you select Earlier than 8.2.1 , make sure that the switch that you want to add is monitored by BNA. BNA is required for monitoring Brocade switches with Fabric OS versions earlier than 8.2.1. Click Next , then specify the connection details and the authentication credentials for the BNA SMI agent that you want to add. You can close the window and continue your work. To change the default probes and performance monitors, keep the window open and update them when the switches are connected.
Cisco	To enable the collection of data with SNMP agents, complete the following steps:
	a. Specify the version of SNMP: v3 (preferred) or v1 .
	b. If you choose to use v3, specify the User name and associated Authentication password. Select an Authentication protocol and an Encryption protocol.
	c. If you choose to use v1 , define the Read community and the Write community .
	d. Identify the Host name or IP address of each switch that you want to add.
	e. Click Next . You can close the window and continue your work. To change the default probes and performance monitors, keep the window open and update them when the switches are connected.

5. Specify the display name and location of the switches and fabrics that are discovered.

The discovered switches are automatically added for monitoring, if they have the same credentials as the switches that you added. A probe is automatically run to collect status and asset information about the resources.

When Brocade switches are added, other switches in the same fabric that do not have the same authentication credentials might also be discovered. These switches are not added for monitoring. They are added to the **Switches** page with a Not Monitored condition.

When the collection of data is complete, you can view status information about the fabric or switch on the **Fabrics** page or **Switches** page.

If you want to monitor multiple switches that have Not Monitored status, you must add them with the correct credentials. To add a list of Not Monitored switches, add the **IP Address** column to the **Switches** page. Then, select Condition as the filter, and Not Monitored as the filter value. Review the switches and gather the IP addresses of the switches that you want to add for monitoring. Then, add the switches in the usual way.

Managing alert conditions and notification settings: Review the alert configuration for the fabric or switch.

If you configured data collection for the switches or fabric, they are automatically added to the default alert policy for their type.

You can change which alert policy manages a switch, and you can also set a switch to not be managed by any policy. For example, to change the alert policy for a switch, go to **Settings** > **Alert Policies** and click the **Resources by Policy** tab. Click **Switches**, select a switch, and then select **Actions** > **Set Policy**.

Restriction: If you don't configure data collection for the switches and fabric, they are not added to the default alert policy.

Related tasks

Sorting lists

Each list of resources, tasks, alerts, or other objects in the GUI has a column or set of columns that determines the order of its rows. You can sort the columns in a list to organize the rows according to your requirements.

Data sources for switches and fabrics

Depending on the type of switch that you want to monitor, you can use Fabric OS REST API, SMI agents, or SNMP agents to monitor switches and fabrics. SMI agents are also known as CIM agents, CIMOM (CIM Object Manager) agents, or SMI-S providers.

The vendors of switches can help you determine which agents to use.

Table 37. Agent types for switch and fabric functions		
Function	Brocade	Cisco
Monitor performance	REST API or SMI agent	SNMP agent
Collect information about switches and switch ports	REST API or SMI agent	SNMP agent
Collect information about topology connectivity	REST API or SMI agent	SNMP agent
Collect information about zoning information and zone control	REST API or SMI agent	SNMP agent
Generate alerts	REST API or SMI agent	SNMP agent

Important: The embedded SMI agent is only available in the Professional Plus and Enterprise editions of Brocade Network Advisor (BNA). BNA is no longer available for purchase, although it will be supported by Broadcom until February 2022. For more information about end of support for BNA, see https://www.broadcom.com/support/fibre-channel-networking/eol.

Adding hypervisors

You can add hypervisors for monitoring by IBM Spectrum Control. Hypervisors include VMware vSphere data sources such as ESX and ESXi, and vCenter Server systems. When you add vSphere data sources, you can then collect data, generate reports, and manage storage that is related to those resources. To add a hypervisor, you specify connection information for the ESX or ESXi data source. You can add multiple hypervisors by specifying connection information for a vCenter Server system.

Adding ESX and ESXi hypervisors

Add ESX and ESXi hypervisors for monitoring by IBM Spectrum Control.

You can add an ESX or ESXi hypervisor by specifying connection information for that hypervisor. For a complete list of hypervisors that you can add, see the ** IBM Spectrum Control interoperability matrix at http://www.ibm.com/support/docview.wss?uid=swg21386446, and go to the *Agents, Servers and Browsers* section.

- 1. In the menu bar, go to **Servers** > **Hypervisors**.
- 2. Click Add Hypervisor.
- 3. Select **ESX/ESXi** and enter connection information about the hypervisor.
- 4. Schedule a probe of the hypervisor.
- 5. Follow the instructions in the GUI to add the hypervisor.

The hypervisor is added for monitoring. A probe is automatically run to collect status and asset information about the hypervisor.

When the collection of data is complete, you can view status information and capacity data about the hypervisor on the **Hypervisors** page.

Managing alert conditions and notification settings: Review the alert configuration for the hypervisor. Each hypervisor is automatically added to the default alert policy for the hypervisor type. You can change which alert policy manages a hypervisor, and you can also set a hypervisor to not be managed by any policy. For example, to change the alert policy for a hypervisor, go to Settings > Alert Policies and click the Resources by Policy tab. Click Hypervisors, select a hypervisor, and then select Actions > Set Policy.

Related tasks

"Adding vCenter Server systems" on page 133

Checking permissions to browse data stores

Determine if the user name that you specified for a VMware data source has permission to browse through the data stores on a hypervisor.

When you add a VMware data source in IBM Spectrum Control, the user name that you specify must have permission to browse through the data stores on VMware. IBM Spectrum Control must browse through the data stores to collect information from the hypervisors. However, the "Read Only" role as defined by VMware does not allow IBM Spectrum Control to browse the data stores. You can use the "Virtual Machine Power User" role if you do not want to use the Administrator role, or you can create a custom role with the required permissions.

To verify that a VMware user is assigned the correct role and privileges to monitor VMware data sources, follow these steps:

- 1. Ensure that the user role has the required VMware datastore permissions by completing the following steps:
 - a) Connect the vSphere Web Client to the VMware data source.
 The data source can be an ESX server, a vCenter Server Appliance, or a vCenter Server.
 - b) From the Inventories view, select Hosts and Clusters.
 - c) Select a host, and click the **Related Objects** tab.
 - d) View the datastores by clicking the **Datastores** tab.
 - e) Right-click a datastore, and select **File Browser**. If you can view the **Files** tab for the datastore, your browse permission is working correctly.
- 2. Determine the role that is assigned to the user by logging in to the vSphere Web Client by using the administrator user ID. From the Administration view, select **Roles**. Verify the role name that is assigned to the user.
- 3. Determine the privileges that are assigned to the role by selecting the user's role and clicking **Privileges**. Expand the privilege groups to view the specific privileges.
- 4. Optional: If you must edit the privileges for the role, select the role and click the **Edit role action** icon. Select privilege groups or expand to select specific privileges.

For more information about VMware user roles, go to the $\underline{\text{VMware documentation center}}$ and search for $\underline{\text{vSphere users and permissions}}$.

Adding vCenter Server systems

For IBM Spectrum Control to monitor multiple hypervisors, you must first add a vCenter Server system.

You can add multiple hypervisors by specifying connection information for a vCenter Server system. For a complete list of hypervisors and vCenter Server systems that you can add, see the Server systems that you can add the Server systems that you can add the Server systems that you

- 1. In the menu bar, go to **Servers** > **Hypervisors**.
- 2. Click Add Hypervisor.
- 3. Select VMware vCenter and enter connection information about the vCenter Server system.
- 4. Schedule a probe for the hypervisors that were discovered.
- 5. Follow the instructions in the wizard to add the vCenter Server system.

After a vCenter Server system is added for monitoring, probes collect status and asset information about its hypervisors.

When the collection of data is complete, you can view status information and capacity data about the hypervisors on the **Hypervisors** page.

Managing alert conditions and notification settings: Review the alert configuration for the hypervisor.

Each hypervisor is automatically added to the default alert policy for the hypervisor type. You can change which alert policy manages a hypervisor, and you can also set a hypervisor to not be managed by any policy.

You can also create an alert policy from the alert definitions and notification settings of a hypervisor. Click **Servers** > **Hypervisors**. Right-click the hypervisor from which you want to create an alert policy, then click **View Alert Definitions**. Click **Create Policy** from the **Policy Actions** menu.

Related tasks

"Checking permissions to browse data stores" on page 133

Determine if the user name that you specified for a VMware data source has permission to browse through the data stores on a hypervisor.

Adding servers

Add servers for monitoring. IBM Spectrum Control creates and updates agentless servers automatically after it probes storage systems and hypervisors. You can also add a server by deploying a Storage Resource agent.

Agentless servers

IBM Spectrum Control creates and updates agentless servers automatically after it probes storage systems and hypervisors.

IBM Spectrum Control creates agentless servers by using information about the following resources:

- · Host connections on storage systems
- · Host bus adapters (HBAs) on fabrics
- · Discovered virtual machines

Depending on your storage environment, IBM Spectrum Control might not model all of the servers in your environment correctly. For example, IBM Spectrum Control might identify six host connections and create an agentless server for each host connection. However, in your environment, these six host connections represent one server computer. In this case, you need to merge the six agentless servers into one agentless server.

Similarly, IBM Spectrum Control might create a single agentless server from a group of host connections that appear to be related. However, in your environment, each of these host connections represents a separate server computer. In this case, you can separate the agentless server into multiple agentless servers, one for each of the servers in your environment.

Related tasks

Tutorial: Reviewing and updating your agentless servers

In this tutorial, you review the agentless servers that are created automatically by IBM Spectrum Control, and update your agentless servers.

Adding servers with Storage Resource agents

You can deploy a Storage Resource agent to a server if you want to view information about the resources that are internal to a server, or about the fabrics that are visible to the server. Internal resources include controllers, disks, file systems, logical volumes, shares, and directories.

To deploy a Storage Resource agent, you must have Administrator privileges.

For each installation of IBM Spectrum Control, you can deploy only one Storage Resource agent on each server. If you attempt to deploy additional Storage Resource agents on a server, the deployments fail.

1. In the menu bar, go to **Servers** > **Servers**.

- 2. Click Deploy Agent.
- 3. Select one of the following methods for deploying Storage Resource agents:
 - Enter information about the server and the Storage Resource agent manually.
 - Import configuration information from a comma-delimited file.
- 4. On the **Deploy Agent** page, configure deployment information for the Storage Resource agents.

If you deploy multiple agents with different operating systems, separate configuration pages are displayed for agents that are deployed on Windows servers and agents that are deployed on UNIX servers.

5. On the **Configure** page, schedule the deployment of the Storage Resource agents.

If you are deploying agents on multiple servers, a time span is calculated during which the agents are deployed. The agents are deployed at regular intervals during the time span to avoid excessive load on the IBM Spectrum Control server.

6. Schedule the time and frequency that probes are run for the servers.

If you add multiple servers, a time span is calculated during which the servers are probed.

7. To add the servers, click **Finish**.

A probe is automatically run for a server after the agent is successfully deployed.

Use **Agent State** on the **Servers** page to monitor the status of the agent deployment.

Managing alert conditions and notification settings: Review the alert configuration for the server.

Each server is automatically added to the default alert policy for the server type. You can change which alert policy manages a server, and you can also set a server to not be managed by any policy.

You can also create an alert policy from the alert definitions and notification settings of a server. Click **Servers** > **Servers**. Right-click the server from which you want to create an alert policy, then click **View Alert Definitions**. Click **Create Policy** from the **Policy Actions** menu.

File List

Add one or more servers with Storage Resource agents by importing the configuration information from a comma-delimited file.

The web-based GUI guides you through the following steps for adding servers:

- Select the input file.
- Configure deployment information.
- Schedule the agent deployment and data collection for the servers.

The comma-delimited file that you use to import the configuration information for servers can contain entries for a single server or multiple servers. Each line in the file represents a server that you want to add. The information about each server must be organized in the following format:

```
host name or IP address,OS type,location,custom tag 1,custom tag 2,custom tag 3
```

where:

- host name or IP address is required for each server entry. An IP address can be in an IPv4 or IPv6 format. A host name or IP address can contain the following characters or symbols:
 - A Z (uppercase characters)
 - a z (lowercase characters)
 - 0 9 (numeric characters)
 - Symbols: . :
- OS type is required and represents the operating system of the server. The OS type for a server must be one of the following values:
 - Windows

- Linux
- AIX
- Solaris
- HP-UX
- location is optional and represents the physical location of the server. The location value can be up to 64 characters in length. If the length exceeds 64 characters, the location value is truncated when the server is added.
- custom tag 1, custom tag 2, and custom tag 3 are optional and represent any additional information that you want to provide about the server. The custom tag values can be up to 64 characters in length. If the length exceeds 64 characters, the custom tag value is truncated when the server is added.

Tip: The custom tags can be displayed on the **Servers** page or can be included as report columns when you generate reports for the server.

If a server has no information for an optional field, you must add a comma separator unless there are no further values for that server.

Example

```
host1, Windows, San Jose, Accounting department host5, Linux, London, Finance department 198.51.100.22, HP-UX,, Computing department 2001:DB8:0:0:0:0:0:0, Windows, Tokyo
```

Tips:

- If the comma-delimited file contains entries for multiple servers, it might take some time to add the servers. To confirm that the servers are added, check the Status column on the **Servers** page.
- To comment out a line, enter a "#" at the beginning of the line. The server on that line is not added when the list is imported.

Example:

```
# host1, Windows, San Jose, Accounting department
```

• If there are syntax problems in the file, none of the servers in the file are added.

Adding rollup servers

Rollup servers gather capacity and status information from IBM Spectrum Control servers in your enterprise. You can add rollup servers for monitoring on the **Settings** > **Rollup Server Connections** page.

IBM Spectrum Control supports master and subordinate rollup servers. Master rollup servers gather capacity and status information about the storage resources that are managed by subordinate rollup servers. You can view this information by enabling IBM Spectrum Control rollup mode. These managed resources include servers, hypervisors, fabrics, switches, and block storage systems.

Configuring rollup servers

To configure rollup servers, you add one or more subordinate servers to a master server. The master server uses probes to collect capacity and status information from the subordinate servers. Master servers must have IBM Spectrum Control 5.3.0 or later installed and running. Subordinate servers must have IBM Spectrum Control 5.3.0 or later installed.

Adding a rollup server

To add and configure subordinate rollup servers, go to **Settings** > **Rollup Server Connections**. You can add one or more subordinate servers to the master (local) server. You can also add a master rollup server as a subordinate server to another master server.

Viewing rollup information

To view information about the resources managed by your rollup servers, you must be in rollup

mode. To enter rollup mode, hover your pointer over the rollup mode icon and select **Enter rollup mode**. You cannot enter rollup mode unless you have added at least one subordinate server to the master (local) server.

Restrictions for rollup servers

When you use rollup mode to view information that was collected by subordinate servers and rolled up to master servers, keep in mind the following limitations, restrictions, and considerations.

Required versions of IBM Spectrum Control

Master rollup servers must have IBM Spectrum Control 5.3.0 or later installed. You can add subordinate servers that run any version of IBM Spectrum Control 5.3.0 or later.

Collecting data after you upgrade

After you upgrade a master rollup server you must run probes of the subordinate servers that were added to the master server prior to the upgrade. Use the Start Probe action to run probes of the subordinate servers. Wait for the first probe (manual or scheduled) to complete, before you work with a new function.

Unmapped Capacity and Mapped Capacity values are slightly different between the subordinate server and the rollup server

Due to differences in the methods that are used to calculate Unmapped Capacity and Mapped Capacity, there might be minor variations between the rollup and subordinate servers. The value that is recorded in the subordinate server resource list panel is the definitive one.

Removing resources

Remove resources that you no longer want to monitor with IBM Spectrum Control.

1. In the menu bar, go to the type of resource that you want to remove.

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Object storage systems	Storage > Object Storage Systems
Servers	Servers > Servers
Hypervisors	Servers > Hypervisors
Switches	Network > Switches
Fabrics	Network > Fabrics
Rollup subordinate servers	Settings > Rollup Server Connections

2. Right-click the resource and click **Remove**.

Tips:

- To remove rollup subordinate servers, click **Actions** > **Remove Server**.
- To remove a Cisco fabric, you must remove all the switches in that fabric. The Cisco fabric is then automatically removed.
- To remove a NetApp storage system configured for block storage, or configured for block storage and file storage, go to **Storage > Block Storage Systems**. If the NetApp storage system is configured only as a filer, then go to **Storage > File Storage Systems**.
- 3. Click **Remove** to confirm that you want to remove the resource.

When you remove a resource, it is no longer monitored by IBM Spectrum Control. All the data that is associated with the resource, including historical data, is removed from the database. All the scheduled tasks and data collection jobs for the resource are canceled.

Collecting data

Determining the data that you want to gather about storage resources is critical to helping you implement a storage management strategy. IBM Spectrum Control provides two different jobs for collecting data about resources.

Table 38. Data collection jobs, resources, and related tasks		
Data collection job	Resources	Related tasks
Probe Use probes to collect status, asset, configuration, and capacity information about resources.	Storage systems Servers Hypervisors Fabrics Switches Rollup servers (capacity and status information only)	 Schedule a probe when you add a resource for monitoring. Modify the schedule for a probe at any time on the list and details pages for a resource. For example, to modify the schedule for a block storage system, go to the Block Storage Systems page or the details page for the block storage system. Use the Rollup server connections page to manage probes for rollup servers.
Performance monitor Use performance monitors to collect metrics that measure the performance of resources.	Block storage systems File storage systems Object storage systems Switches	 Schedule a performance monitor when you add a resource for monitoring. Modify the schedule for a performance monitor at any time on the Performance page, and on the list and details pages for the resource.

Related tasks

Viewing capacity alerts and violations

View the alerts that were triggered when the capacity of a resource changes and reaches a threshold. For example, you can view alerts that are generated when the measured value of a capacity metric for a pool meets the conditions for generating an alert.

Defining alert definitions for general attributes and capacity changes

You can define alerts that are triggered when the attributes or capacity of a resource changes. Attributes represent the key properties and configuration of a resource, such as status, versions, removals, discoveries, and data collection status. Capacity represents storage statistics such as available capacity, used capacity, drive capacity, reserved capacity, and more.

Collecting asset and status data by using IBM Spectrum Control probes

Use IBM Spectrum Control probes to collect asset, status, and storage data about monitored resources in your environment. Probes are also used to discover information about new or removed disks and file systems. You can view information about probes on the resource pages in IBM Spectrum Control, for example, the **Block Storage Systems** page, or the Block Storage Systems details page.

Creating probes in IBM Spectrum Control

Use probes to collect asset, status, and storage data about resources that are monitored by IBM Spectrum Control. You can create probes when you add resources for monitoring or after you add the resources for monitoring.

You can create probes for one or more resources at the same time. You can also create an automated or manual probe. Schedule an automated probe to manage your probe jobs efficiently and to minimize the load on the IBM Spectrum Control server. If you want to run a probe at a specific time, schedule a manual probe.

1. To schedule a probe for a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Object storage systems	Storage > Object Storage Systems
Servers	Servers > Servers
Hypervisors	Servers > Hypervisors
Switches	Network > Switches
Fabrics	Network > Fabrics
Rollup subordinate servers	Settings > Rollup server connections

Tip: For switches, probes also collect asset and topology information about the fabrics that include those switches.

- 2. Locate the resources that you want to probe.
- 3. Select the resource rows and click **Actions** > **Data Collection** > **Schedule**. For Rollup subordinate servers, select **Actions** > **Start probe** for the server.
- 4. On the **Data Collection Schedule** window, set the probe to **Enabled**.
- 5. From the following table, choose the type of probe that you want. Make your selections for the data collection, and save your selections.

Type of probe	Selections for probe	
Automated probe	Select the frequency of the data collection. Click Save .	
Manual probe	Select the time and the frequency of the data collection. Click Save .	

For automated probes, the start time is determined by an algorithm that minimizes the number of concurrent probe jobs to avoid excessive load on the IBM Spectrum Control server. Automated probe jobs are scheduled by IBM Spectrum Control within the time interval that you define in the **Automated Probe Run Window**.

Tips for manually scheduling probes:

• When you schedule probes to collect data about resources, schedule the probes to run at different times. By scheduling data collection so that no two probes run at the same time, you can help improve the overall performance of the product and ensure that data is collected successfully.

• If you schedule manual probes for multiple resources, a time span is calculated during which the resources are probed. The resources are probed at different times during the time span to avoid excessive load on the IBM Spectrum Control server.

Probe schedules are created for the resources that are selected. The probe jobs run according to the modified schedule. If you disabled the probes, IBM Spectrum Control stops collecting asset, status, and capacity data for the monitored resources.

Note: Status and asset information about logical switches is collected when the parent switch is probed.

You select three storage systems, storage_system_a, storage_system_b, and storage_system_c. A probe schedule was defined for storage_system_c, but not for storage_system_a and storage_system_b. You set the time and frequency fields, and set the probe to **Enabled**. When you click **Save**, probe schedules are created for storage_system_a and storage_system_b. The probe schedule that was defined for storage_system_c remains unchanged.

How automated probes are scheduled

When you schedule probes, you can manually specify when they run, or you can have IBM Spectrum Control automatically schedule them. Automated probe scheduling helps remove the complexity of manually planning when probes are run and ensures that server resources (CPU, Memory, and so on) required for data collection are distributed over time. Distributing the use of system resources over time helps avoid excessive load on IBM Spectrum Control at any one time.

To schedule probes automatically, an algorithm is used to determine the best times for when the probes are run. To determine this time, the algorithm considers the existing schedules of probes for all types of resources. For example, when an automated probe of a storage system is scheduled, IBM Spectrum Control attempts to minimize the number of probes that are running concurrently by analyzing the existing probes of switches, servers, hypervisors, and other storage systems.

IBM Spectrum Control also analyzes the frequency, start time, and expected duration of existing probe schedules. This analysis includes probes that were scheduled automatically or manually. After the analysis, IBM Spectrum Control creates a unique probe schedule for each day of the week.

Keep in mind the following information about automatic probe scheduling:

- When you schedule a new probe during the week, the schedule of existing probes does not change. Existing probes are considered when the schedule of the new probe is determined, but they continue to run according to the previously defined schedule for that week.
- When a probe is scheduled to run on a frequency of every *X* days where *X* > 1, IBM Spectrum Control determines the days for that frequency based on a number of conditions. Some of the conditions include whether the probe was run previously as a scheduled probe and how much time is available until the end of the current probe window.
- Automated probes can be rescheduled by IBM Spectrum Control under the following conditions:
 - When you change the time range for the automated probe schedule, IBM Spectrum Control reschedules all the automated probes to optimize when probes are run during that new time range.
 - On Saturday night (in the time zone of the IBM Spectrum Control server) of every week, all automated probes are rescheduled. During the rescheduling process, the algorithm is again applied to determine the best times to run automated probes on each day of the upcoming week. Because all automated probes are rescheduled weekly, IBM Spectrum Control ensures that the average number of concurrent probes is kept minimized for your environment on an ongoing basis.

Verifying that a probe is running for a resource

Verify that a probe is running for a resource to ensure that IBM Spectrum Control is collecting asset and status information about the resource.

To verify that a probe is running for a resource, use the resource list page in the web-based GUI.

- 1. From the menu bar, go to a resource list page for a resource type.

 For example, to verify that a probe is running for a storage system, go to **Storage** and select the type of storage system you want to monitor.
- 2. Locate the resource that you want to check.

- 3. View the Probe Status column.
 - A value of Successful or Running indicates that the probe is collecting asset and status information about the resource. The Last Successful Probe column shows the most recent date and time when the probe successfully collected data.
- 4. Optional: If the Probe Status column shows a status of Failed or Warning, right-click the resource and click **Data Collection** > **Open Probe Logs** to view the messages in the log file.
- 5. Optional: If the Probe Status column shows a status of Never probed, right-click the resource and click **Data Collection** > **Schedule** to schedule a probe for the resource.

Modifying probes

Probes are data collection jobs that collect status and asset information about monitored resources in your environment. To modify a probe in IBM Spectrum Control, use the resource list page for the monitored resource in the web-based GUI.

You can select one or more devices and modify their schedules.

1. To modify a probe for a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Object storage systems	Storage > Object Storage Systems
Servers	Servers > Servers
Hypervisors	Servers > Hypervisors
Switches	Network > Switches
Fabrics	Network > Fabrics

Tip: For switches, probes also collect asset and topology information about the fabrics that include those switches.

- 2. Locate the resources that you want to probe.
- 3. Select the resource rows and click **Actions** > **Data Collection** > **Schedule**.
- 4. On the Data Collection Schedule window, modify the schedule for the probes, and click Save.

For automated probes, the start time is determined by an algorithm that minimizes the number of concurrent probe jobs to avoid excessive load on the IBM Spectrum Control server. Automated probe jobs are scheduled by IBM Spectrum Control within the time interval that you defined in the **Automated Probe Run Window**.

If you schedule manual probes for multiple resources, a time span is calculated during which the resources are probed. The resources are probed at different times during the time span to avoid excessive load on the IBM Spectrum Control server.

The following rules determine the settings for the probe fields when you open the **Data Collection Schedule** window:

- The Enable Probe check box is checked by default. Uncheck this box to disable the probe operation for all selected resources.
- If all selected resources are configured with automated probe scheduling, then Automatically is selected; if all selected resources are configured with manual probe scheduling, then Manually is selected. If some resources are configured with automatic probe scheduling and some resources are configured with manual probe scheduling, then neither Automatically nor Manually is selected.
- Fields that are blank indicate that more than one value (Automatically or Manually) is used among the selected resources. If the field contains a red asterisk (*), you must set a value. If the field is blank, the existing values are preserved unless you select a new value to be applied to all of the selected resources.

The probe jobs run according to the modified schedule. If you disabled the probes, IBM Spectrum Control stops collecting asset and status data for the monitored resources.

The changes that you make to the probe schedule are applied to all the resources that you select. For example, you select two servers with different probe frequency values. Probes are disabled for both servers. When you open the **Data Collection Schedule** window, the probe status field is set to **Disabled** and the frequency field is blank. You set the status field to **Enabled** and the frequency field to Every 2 days. When you click **Save**, the probe is enabled with the new frequency value for the two servers that you selected

Configuring alerts for probes

Configure the alert notifications that are generated if a probe fails to run.

You can define alerts for a resource if you want to be notified when the following conditions occur:

- The status of its probe is error or warning. An error status occurs when a probe did not complete and no data was collected about a resource. A warning status occurs when a probe completes, but might not have collected a complete set of data.
- A specified amount of time has passed since a probe collected data about a resource. You can use this type of alert to be notified when up-to-date configuration and status data is not being collected about a resource and its existing data might be stale.

To define an alert for the probe that collects data about a resource, complete the following steps:

1. To select a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Object storage systems	Storage > Object Storage Systems
Servers	Servers > Servers
Hypervisors	Servers > Hypervisors
Switches	Network > Switches
Fabrics	Network > Fabrics

- 2. Right-click a resource in the list and click **Edit Alert Definitions**.
- 3. To enable the **Probe Status** alert for an attribute, click the corresponding switch icon and select a status to be alerted on.

You can use this alert to be notified when a probe is only collecting partial data about a resource or not collecting any data.

4. To enable the **Last Successful Probe** alert for an attribute, click the corresponding switch icon and specify how much time can pass before an alert is generated.

You can use this alert to be notified when up-to-date configuration and status data is not being collected about a resource and its existing data might be stale.

5. Click **Save**.

Starting probes

Use the **Start Probe** action to immediately collect data about resources.

If a probe job is not defined for the resource, use the **Schedule** action to create a probe.

You can start probes for multiple resources at the same time.

Tip: When you start a probe for a switch, the probe gathers statistics about the fabric that the switch is a part of.

1. From the menu bar, go to the resource list page for a resource type.

For example, to start a probe for a storage system, go to **Storage** and select the type of storage system you want to monitor.

- 2. Locate the resources that you want to probe.
- 3. Select the resource rows and click Actions > Data Collection > Start Probe.

The **Start Probe** action is available for the resources that you select if at least one of the resources meets the following criteria:

- A probe job is defined for the resource.
- The probe job is enabled.
- · The probe job is not currently running.

The probe status changes to Running.

Probes are started for all the resources that meet the criteria. For example, you select two storage systems, Storage_System_a and Storage_System_b. Probes are enabled for both of the storage systems and a probe is running for Storage_System_a. The probe is started for Storage_System_b when you select **Start Probe**. Storage_System_a is not affected by the action because a probe is already running for the resource.

Viewing probes for a specific resource

You can view information for specific resources on the resource details pages. For example, use the Storage Systems details page to view the date and time when data was last collected.

- 1. In the menu bar in the web-based GUI, go to a type of top-level resource.

 For example, if you want to view the resource information for a storage system, go to **Storage** and select the type of storage system you want to monitor.
- 2. Locate the resource that you want to view information for.
- 3. Right-click the resource row and select View Details.
- 4. In the General section of the resource details page, click Data Collection.

The Data Collection pane is displayed, which contains information about the data collection jobs that are defined for the resource.

Note: To view information about rollup resources, you must be in *rollup mode*. To enter rollup mode,

move the mouse pointer over the rollup icon in the menu bar and select **Enter rollup mode**.

Viewing probe logs

A probe log file is created for each run of a probe on a storage resource. Probe logs include detailed information about the status, actions, and progress of a probe.

- In the menu bar in the web-based GUI, go to a resource list page for a resource type.
 For example, to view the probe logs for a storage system, go to **Storage** and select the type of storage system you want to monitor
- 2. Locate the resource that you want to view.
- 3. Right-click the resource row, and select **Data Collection** > **Open Probe Logs**.

Restriction: If you select multiple rows, the **Open Probe Logs** action is not available.

The **Logs** page opens and shows log entries for the most recent run of the probe.

- 4. Optional: On the Logs page, to view a previous probe log, select a log from the Select a log list.
- 5. Optional: To view only the log entries that have a Warning or Error status, select an option from the **Show all** list.

You can choose to view only entries that have the following statuses:

- Only error entries
- A Only warning entries

- 8 Error and warning entries
- 6. Optional: To view an explanation of the message that is associated with a log entry, click the link in the **ID** column.

You can view the following information on the **Logs** page:

- The overall status of a probe. The icon that is shown in the Select a log list represents the most critical status that was generated by an action in the job run.
- The status for each probe action.
- The date and time when an action was completed. The date, time, and time zone of an action is shown in the Date and Time column.
- The description of an action.

Collecting performance data by using IBM Spectrum Control performance monitors

Before you can monitor the performance of the resources in your environment, you must collect data about those resources. IBM Spectrum Control uses performance monitors to collect metrics for measuring the performance of storage systems and switches. Use the **Performance Monitors** page to view and manage performance monitors.

You must schedule and run performance monitors before you can complete other tasks, such as optimizing storage tiering, balancing pools, and running performance reports.

Before you can view performance information for resources, you must complete the following tasks:

- Add the storage system or switch for monitoring by IBM Spectrum Control.
- Schedule a performance monitor to collect performance data about the resource.
- Define a performance alert so that you are notified when the performance of a resource might represent a potential problem. For example, you can define an alert threshold that notifies you when the total I/O port rate for a storage system falls outside a specified range.

The performance information that is gathered includes metrics that measure the performance of volumes, ports, and disks. You can view performance metrics on resource list and detail pages for switches and storage systems. For example, you can view performance metrics that are related to switches in the **Performance** tab on the **Switches** page.

Creating performance monitors in IBM Spectrum Control

Performance monitors are data collection jobs that collect metrics for measuring the performance of storage systems and switches. You can create performance monitors in IBM Spectrum Control when you add resources for monitoring or at a later time.

To create performance monitors, you must have Administrator privileges.

To create performance monitors for resources in your environment, you can use the list and detail pages for the resources. For example, to create performance monitors for switches, use the **Switches** page. You can create performance monitors for multiple resources at the same time.

Restrictions:

- You can define only one performance monitor for a resource.
- You can create performance monitors for the following resources only:
 - Storage systems
 - Switches that are managed by a CIM agent
- 1. From the menu bar, go to a resource list page for a resource type.

 For example, to create performance monitors for storage systems, go to **Storage** and select the type of storage system you want to monitor.
- 2. Locate the resources that you want to monitor.
- 3. Select the resource rows and click **Actions** > **Data Collection** > **Schedule**.

- 4. On the Data Collection Schedule window, set the performance monitors to Enabled.
 - You can set a performance monitor to **Enabled** only if the probe is set to **Enabled**.
- 5. Select the interval for the performance monitors, and click **Save**.

The interval represents the number of minutes over which samples of performance data are averaged. If you are creating performance monitors for multiple resources, the interval list displays the performance monitor intervals that are common to all resources.

For example, you select two storage systems. Storage_System_a specifies performance monitor intervals of 5 minutes, 10 minutes, and 15 minutes. Storage_System_b specifies intervals of 10 minutes, 15 minutes, and 20 minutes. The interval list displays intervals of 10 minutes and 15 minutes.

Restrictions:

Note the following restrictions when you create performance monitors:

- a. If you specify a 1-minute interval for performance monitors, data collection samples might be inconsistent lengths. The inconsistent lengths occur because it might take longer than 1 minute to collect the data on some systems. The inconsistent lengths can cause the data in the historical line chart to be misleading. This problem can also occur for the 5-minute summary chart when a 1-minute interval was specified for data collection.
 - Check the messages in the job log to determine the length of time that data collection takes for a resource. If the data collection takes longer than 1 minute, consider changing the interval to 5 minutes. You can use a 1-minute interval when you want to troubleshoot a resource.
- b. If you specify a 1-minute interval, the amount of data that is stored in the database repository increases significantly. IBM Spectrum Control stores only 7 days of sample data that is collected at 1-minute intervals.

Performance monitors are created for the resources that are selected, and that meet the criteria. For example, you select three storage systems, Storage_System_a, Storage_System_b, and Storage_System_c. You enable the performance monitors and set the interval field. When you click **Save**, performance monitors are created for Storage_System_a and Storage_System_b. Storage_System_c is not affected by the action because you cannot configure performance monitors for that storage system.

If a successful probe run is completed for a resource, the performance monitor runs according to the defined interval.

To check the progress of a performance monitor, you can complete the following actions:

- From the resource list page, right-click a resource row, and select Data Collection > Open
 Performance Monitor Logs. You can view detailed informational, warning, and error messages that are related to the performance monitor. Use this information to troubleshoot any errors that might occur when the performance monitor runs.
- View the Performance Monitor Status column on the resource list page.

Verifying that a performance monitor is running for a resource

Verify that a performance monitor is running for a resource to ensure that IBM Spectrum Control is collecting performance metrics about the resource. Use performance monitors to collect metrics about the performance of storage systems and switches.

To verify that a performance monitor is running for a resource, use the resource list page in the webbased GUI.

- 1. From the menu bar, go to a resource list page for a resource type.

 For example, to verify that a performance monitor is running for a storage system, go to **Storage** and select the type of storage system you want to monitor.
- 2. Locate the resource that you want to check.
- 3. View the Performance Monitor Status column.

A value of Completed or Running indicates that the performance monitor is collecting performance metrics about the resource. The Last Successful Monitor column shows the most recent date and time when the performance monitor successfully collected data.

- 4. Optional: If the Performance Monitor Status column shows a status of Running with problems, Completed with warnings, Failed, or Canceled, you can view the logs to troubleshoot the error or warning messages. To view the performance monitor logs, right-click the resource row and click Data Collection > Open Performance Monitor Logs.
- 5. Optional: If the Performance Monitor Status column shows a status of Disabled, to enable the performance monitor for the resource, right-click the resource row and click **Data Collection** > **Schedule**. On the **Data Collection Schedule** page, select **Enabled**.

Tip: You can set the performance monitor to **Enabled** only if the probe is set to **Enabled**. If a probe is set to **Disabled**, the performance monitor cannot run.

Modifying performance monitors

To modify all of the performance monitors that are used by IBM Spectrum Control, use the **Performance Monitors** page in the web-based GUI. To modify performance monitors for specific resources, use the resource list and resource detail pages.

To modify performance monitors, you must have Administrator privileges.

You can modify the schedule for multiple performance monitors at the same time.

You can enable or disable performance monitors and modify the interval. The interval represents the number of minutes over which samples of performance data are averaged.

- In the menu bar in the web-based GUI, go to Home > Performance Monitors.
 Detailed information about all performance monitors is shown in the Performance Monitors tab.
- 2. Locate the performance monitors that you want to modify.
- 3. Select the performance monitor rows and click **Actions** > **Schedule**.
- 4. On the **Data Collection Schedule** window, modify the schedule for the performance monitors.

If you select multiple performance monitors, the following rules are used to determine the field settings when you open the **Data Collection Schedule** window:

- If all of the performance monitors that you select are enabled, the **Enabled** value is displayed. If all of the performance monitors are disabled, the **Disabled** value is displayed. Otherwise, the state field for the performance monitors is blank.
- If the performance monitors have different interval values, the interval field is blank. If the performance monitors have the same interval value, the value is displayed.

Intervals that are common to all performance monitors that you select are displayed in the interval list. For example, you select performance monitors for two storage systems. Storage_System_a specifies performance monitor intervals of 5 minutes, 10 minutes, and 15 minutes. Storage_System_b specifies intervals of 10 minutes, 15 minutes, and 20 minutes. The interval list displays intervals of 10 minutes and 15 minutes.

5. Click Save.

The changes are applied to all the performance monitors that are selected. The performance monitors run according to the modified schedule. If you disabled a performance monitor, IBM Spectrum Control stops collecting performance metrics for the storage system or switch.

To check the progress of the performance monitors, you can complete the following actions on the **Performance Monitors** page:

- View information such as the status of the performance monitor and the most recent date and time when performance data was collected about a resource.
- Right-click the performance monitor row, and select **Open Logs**. You can view detailed informational, warning, and error messages that are related to the performance monitor job. Use this information to troubleshoot any errors that might occur when the job runs.

Configuring alerts for performance monitors

Configure the alert notifications that are generated if a performance monitor fails to run.

You can define alerts for a storage system or switch if you want to be notified when the following conditions occur:

- The status of its performance monitor is error or warning. An error status occurs when a performance monitor did not complete and no performance data was collected about a resource. A warning status occurs when a performance monitor completes, but might not have collected a complete set of performance data.
- A specified amount of time has passed since a performance monitor collected data about a resource. You can use this type of alert to be notified when up-to-date performance data is not being collected about a resource and its existing data might be stale.

To define an alert for the probe that collects data about a resource, complete the following steps:

1. To select a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Switches	Network > Switches

- 2. Right-click a resource in the list and click **Edit Alert Definitions**.
- 3. To enable the **Performance Monitor Status** alert for an attribute, click the corresponding switch icon and select a status to be alerted on.

You can use this alert to be notified when a performance monitor is only collecting partial data about a resource or not collecting any data.

4. To enable the **Last Successful Monitor** alert for an attribute, click the corresponding switch icon and specify how much time can pass before an alert is generated.

You can use this alert to be notified when up-to-date performance data is not being collected about a resource and its existing data might be stale.

5. Click Save.

Starting and stopping performance monitors

To start or stop performance monitors for storage systems and switches, use the **Performance Monitors** page in the GUI.

To start or stop performance monitors, you must have Administrator privileges.

You can start or stop performance monitors for multiple resources at the same time.

Restriction: A successful probe must be completed on the resource before a performance monitor can start. If a successful probe is not run, you cannot start the performance monitor for the resource.

- 1. In the menu bar in the GUI, go to **Home** > **Performance Monitors**.
- 2. Click the **Performance Monitors** tab.
- 3. Locate the performance monitors that you want to start or stop.
- 4. Select the performance monitor rows and click **Actions** > **Start** or **Actions** > **Stop**.

Performance monitors are started or stopped for all the resources that meet the criteria. For example, you select two storage systems, Storage_System_a and Storage_System_b. Performance monitors are enabled for the storage systems and both resources are probed. A performance monitor is running for Storage_System_a but not for Storage_System_b. The performance monitor is started for Storage_System_b when you select **Start**. Storage_System_a is not affected by the action because a performance monitor is already running for the resource.

Starting or stopping a performance monitor has the following effects:

Start

Starting a performance monitor starts the immediate collection of data from the monitored resource. The status of the performance monitor changes to Running.

Stop

Stopping a performance monitor stops the collection of performance data from the resource. If a performance monitor is running, it stops. Performance data is not collected from the resource until you restart the performance monitor. The status of the performance monitor changes to Canceled.

To check the progress of the performance monitor, you can complete the following actions on the **Performance Monitors** page:

- View information such as the status of the performance monitor and the percentage of data collections that succeeded during the most recent 24 hours when the performance monitor was active.
- Right-click the performance monitor row, and select **Open Logs**. You can view detailed informational, warning, and error messages that are related to the performance monitor. Use this information to troubleshoot any errors that might occur when the performance monitor runs.

Viewing performance monitors

To view all of the performance monitors that IBM Spectrum Control uses to collect metrics for measuring the performance of resources, use the **Performance Monitors** page. To view the performance monitors for specific resources, use resource list and resource details pages.

Viewing all performance monitors

To view all of the performance monitors that IBM Spectrum Control uses to collect metrics for measuring the performance of storage systems and switches, use the **Performance Monitors** page.

The **Performance Monitors** page displays one row for each storage system or switch that IBM Spectrum Control is monitoring. If a performance monitor is not scheduled for a resource, the status column is set to Disabled for the performance monitor row.

Before IBM Spectrum Control can collect performance metrics for a resource, the resource must fulfill the following conditions:

- A probe must be run for the resource.
- For switches, the switch must be managed by a CIM agent.
- 1. In the menu bar, go to **Home** > **Performance Monitors**.
 - On the **Performance Monitors** tab, the information about performance monitors is organized into columns. These columns include the status of performance monitors, the names of the monitored resources, and the most recent dates and times when data was successfully collected.
- 2. Optional: View the status icons at the top of the page to view a summary of performance monitor statuses.

This summary includes the number of performance monitors that have a Running, Running with problems, Failed, or Not Running status.

Viewing performance monitors for specific resources

To view the performance monitors for specific storage systems and switches, use the resource detail pages in the web-based GUI. For example, use the **Storage Systems** detail page to view performance monitor information for a specific storage system.

If performance monitor information is not displayed on the resource detail page, this indicates that the resource does not support the collection of performance metrics by IBM Spectrum Control.

- 1. In the menu bar in the web-based GUI, go to a type of top-level resource.

 For example, if you want to view the performance monitor for a storage system, go to **Storage** and select the type of storage system you want to monitor.
- 2. Right-click a resource in the list, and select **View Details**.

The resource detail page is displayed.

In the **General** section, the number next to **Data Collection** represents the number of data collection jobs that are associated with the resource. For example, Data Collection (2) indicates that you can view a performance monitor and a probe for the resource.

The icon next to the **Data Collection** link represents the most critical status of the data collection jobs that are associated with the resource. For example, if a job failed, the following icon is shown:

Data Collection (2)



3. In the General section, click Data Collection to view information about the performance monitor for the switch or storage system.

On the **Data Collection** pane, you can view information such as the status of the most recent run of a performance monitor, the performance monitor interval, and the most recent data and time that data was successfully collected.

Tip: You can use the columns on the resource list pages to view information about the performance monitors that are defined for storage systems or switches. For example, use the Storage Systems page to view performance monitor information for storage systems.

Viewing performance monitors logs

Use performance monitor logs to view detailed information about the status, actions, and progress of a performance monitor. You can use this information to troubleshoot any errors that might occur when a performance monitor is running.

Performance monitor logs contain informational, warning, and error messages that are related to each action that is taken during the processing of a performance monitor.

- 1. In the menu bar in the web-based GUI, go to **Home > Performance Monitors**.
- 2. In the **Performance Monitors** tab, locate the performance monitor that you want to view.
- 3. Right-click the performance monitor row, and select **Open Logs**.

Restriction: If you select multiple rows, the **Open Logs** action is not available.

The **Logs** page opens and shows log entries for the most recent performance monitor log. A new log file is created when you stop and restart a performance monitor.

- 4. Optional: On the Logs page, to view a previous performance monitor log, select a log from the Select a log list.
- 5. Optional: To view only the log entries that have a Warning or Error status, select an option from the Show all list.

You can choose to view only entries that have the following statuses:



• Only error entries



• Only warning entries





• S Error and warning entries

6. Optional: To view an explanation of the message that is associated with a log entry, click the link in the **ID** column.

The information on the **Logs** page is automatically updated every 30 seconds. New entries are added to the end of a log. You can view the following information on the **Logs** page:

- The overall status of a performance monitor. The icon that is shown in the **Select a log** list represents the most critical status that was generated by an action in the job run.
- The status for each performance monitor action.
- The date and time when an action was completed. The date, time, and time zone of an action is shown in the Date and Time column.

- The ID of the message that is associated with an action. You can click the value in the ID column to view more information about a message.
- The description of an action.

Tip: You can use the navigation and search functions of the web browser to locate information on the page.

Collecting information about shares on storage systems

To collect information about shares in file systems, you must modify the TPCD.config file. When you provision shares, you can then view such information as the name of the share, the name of the file system, and the protocols that are used to share files.

Set the **saveNonRoot** parameter to **saveNonRoot=1** in the TPCD. config file to view the information that is collected about shares in the GUI.

1. Open the TPCD.config file.

For Windows operating systems:

By default, the TPCD.config file is in the *installation_dir*\data\config\ directory.

For Linux or UNIX operating systems:

By default, the TPCD.config file is in the *installation_dir*/data/config/ directory.

Where installation_dir is where IBM Spectrum Control is installed.

2. In the [server] section of the TPCD.config file, change the value of the **saveNonRoot** parameter to 1 as shown in the following example:

```
saveNonRoot=1
```

- 3. Save the TPCD. config file.
- 4. Restart the Data server.

When you provision new shares, you can view information about the shares on the **Provisioning** task page in the GUI. On the **Shares** page, you can view more information about shares, for example, the servers that have access to the share.

Collecting information about the sizes of snapshots in IBM Spectrum Scale

By default, IBM Spectrum Control does not collect information about the size of GPFS snapshots in IBM Spectrum Scale. To view that information, you must enable IBM Spectrum Control to collect information about the size of snapshots.

Restriction: When you enable data to be collected about the size of GPFS snapshots, the collection might put a significant load on the IBM Spectrum Scale cluster during the probe. This load might affect other applications that use the same cluster.

1. Run the following command to set the *Probe.GetGPFSSnapshotSize* property for the Device server to true:

```
tpctool setdscfg -user user\_ID -pwd password -url localhost:9550 -property Probe.GetGPFSSnapshotSize true
```

Where *user_ID* is an IBM Spectrum Control user ID and *password* is the password for the IBM Spectrum Control user ID.

2. Restart the IBM Spectrum Control Device server and the web server.

Alerting

Specify conditions that trigger alerts and the actions to take when those alerts are triggered, such as notify an email address. Use alert policies to define those alert conditions and notification settings for a group of resources.

Alerting functions examine the attributes, capacity, and performance of resources. If the conditions that are defined for alerts are met, the actions that are specified for the alert are taken. Typically, the actions include sending a notification. For example, if the status of a SAN Volume Controller storage system changes to Error, an alert is displayed in the **Alerts** page in the GUI, and an email might be sent to a storage administrator.

You can manage alerts in your storage environment in the following ways:

- Use alert policies to manage the alert definitions and notification settings that apply to different sets of resources. For example, you can use one alert policy for the servers in your test environment, and another for the servers in your production environment. Here are some important points about alert policies:
 - Alert policies manage one type of resource only. For example, if you have SAN Volume Controller and FlashSystem 900 storage systems in your storage environment, you cannot have both types of resource in one alert policy.
 - A resource can be managed by only one alert policy.
 - When you add a resource to be monitored by IBM Spectrum Control, it is added to a default alert policy automatically. However, default alert policies are not provided for agentless servers.
 - If a resource is managed by a policy, the resource cannot have alert definitions and notification settings that are independent of the policy. The alert definitions and notification settings that apply to the resource come from the policy.
 - It is not a requirement for resources to be managed by an alert policy. A resource can have its own alert definitions and notification settings, independent of an alert policy.
 - When you add a resource to be monitored by IBM Spectrum Control, it is added to a default alert policy automatically.
 - Default policies with alerts already configured are available. You can create copies of the default policies and assign resources to the new policies. Your alerts are configured with the default settings.
 - If you add a resource to a policy, any existing alert definitions for the resource are replaced by the
 alert definitions in that policy. You can't restore the original alert definitions for a resource after you
 move it into the policy.
 - To store a copy of the original alert definitions for a resource before you move it into a policy, create a policy based on the alert configuration of the resource. Then, move the resource to another policy. If you want to reapply the original alert definitions later, you can add the resource to the policy with the original alert configuration.
- Define alert conditions and notification settings for individual resources. It is not a requirement for resources to be managed by an alert policy. A resource can have its own alert definitions and notification settings, independent of an alert policy.
- Define alerts and notification settings for applications and general groups. Use applications or general
 groups to manage alerts for groups of resource components such as volumes or pools. For example, you
 might want to define alerts on the response time for volumes in an application, depending on the
 response time requirements of the application. In this case, it is not useful to configure volume
 response time thresholds for the entire storage system because the storage system might serve many
 different applications with different needs.

Related tasks

Adding resources

You must add resources to IBM Spectrum Control before you can collect data, generate reports, and manage storage that is related to those resources.

Viewing capacity alerts and violations

View the alerts that were triggered when the capacity of a resource changes and reaches a threshold. For example, you can view alerts that are generated when the measured value of a capacity metric for a pool meets the conditions for generating an alert.

Defining alert definitions for general attributes and capacity changes

You can define alerts that are triggered when the attributes or capacity of a resource changes. Attributes represent the key properties and configuration of a resource, such as status, versions, removals, discoveries, and data collection status. Capacity represents storage statistics such as available capacity, used capacity, drive capacity, reserved capacity, and more.

Related reference

Triggering conditions for alerts

Define alerts so that IBM Spectrum Control automatically notifies you when certain conditions or events are detected on monitored resources. Such conditions are the *triggering conditions* for the alert. The specific conditions that can trigger alerts depend on the type of resource that is being monitored.

How alerts work

Alerting functions examine the attributes, capacity, and performance of resources. If the conditions that are defined for alerts are met, the actions that are specified for the alert are taken. Typically, the actions include sending a notification. For example, if the status of a SAN Volume Controller storage system changes to Error, an alert is displayed in the **Alerts** page in the GUI, and an email might be sent to a storage administrator.

Triggering conditions for alerts

Event processing

Conditions that generate alerts are detected during data collection and event processing. By default, probes collect storage system data once every 24 hours. For some storage systems such as IBM Spectrum Accelerate and the XIV, events are polled every minute from the resource. For IBM Spectrum Scale, status change events are polled frequently, typically within minutes. For other resources, events are subscription-based, where the resource itself or a data source such as a CIM agent sends the events to IBM Spectrum Control when conditions change on the resource.

Examples of storage systems that use subscription-based event processing include SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, FlashSystem V840, and FlashSystem V9000. For these storage systems, a probe is automatically run when many events are received from the storage system in a short time period. To avoid performance bottlenecks, probes are run only every 20 minutes.

Determining which type of alert to use

To determine whether to define alerts in alert policies, for individual resources, or for the set of resources that are included in an application or general group, follow these guidelines:

Which type of alerts to use?	Scenario
Alerts defined in alert policies	You want to manage alert conditions and notification settings for a group of resources of the same type. For example, if you have several SAN Volume Controller storage systems in your environment, you can create an alert policy so that the alert definitions are the same for all of the SAN Volume Controller systems.
	If you have some SAN Volume Controller systems in a test environment, and some in a production environment, you can use one alert policy for the test environment, and another for the production environment.

Which type of alerts to use?	Scenario
Resource alerts	You want to receive alert notifications about the performance of a specific resource, or its internal resources. For example, for a storage system, you can alert on the performance of the storage system itself, and on the performance of its volumes and pools.
	If you define an alert for a resource, for example, a performance alert for the ports on a storage system, the alert threshold value applies to all of the ports on the storage system. You cannot apply different alert thresholds to internal resources of the same type on a resource.
Application alerts	Use application alerts in the following scenarios:
	You want to receive alert notifications for all the resources of a certain type in an application. For example, if your application uses multiple storage systems, you can define the storage system alerts once for the application and the alerts apply to all the storage systems. If you later add more storage systems to the application, the existing application alerts apply to those storage systems also.
	You want to apply different thresholds to internal resources of the same type on a storage system. For example, you have production applications and test applications that use volumes on a SAN Volume Controller. The production applications require response times of 6 milliseconds or less while the test applications can tolerate response times up to 30 milliseconds. You can use application alerts to set separate response time thresholds for volumes used by the different applications, depending on the needs of that application.
General group alerts	Use general group alerts in the following scenarios:
	You want to receive alert notifications about changes for a subset of the resources of a particular type. For example, you can detect when the ports that are used for replication on your SAN Volume Controller have insufficient buffer-to-buffer credit. Alert notifications are not generated for ports that are not used for replication.
	 You want to receive alert notifications about changes for a group of resources that are logically related. You can group all the storage systems at a specific location or all the servers that use a particular operating system. For example, you can receive alert notifications when the used capacity of any of your Linux servers exceeds 80%.

Tip: If a resource is in both an alert policy and a general group, the alert definitions for both the policy and the group are applied.

For detailed information about how to define alerts and view alert notifications, see the <u>IBM Knowledge</u> Center for IBM Storage Insights.

Viewing and administering alerts

View and administer the alert that were detected on monitored resources, applications, and general groups.

Table 39. Viewing and administering alerts	
Actions	Steps
Remove alerts.	 For resources, go to the resource list page for the resource. For applications, go to Groups > Applications. For general groups, go to Groups > General Groups.
	Right-click a resource, application, or general group and click View Details.
	3. Click Alerts in the General section.
	4. Right-click one or more alerts and click Remove alerts .
	5. Click Remove .
	To remove all alerts go to Home > Alerts , and then click Remove all alerts in the Actions menu.
Acknowledge alerts.	 For resources, go to the resource list page for the resource. For applications, go to Groups > Applications. For general groups, go to Groups > General Groups.
	Right-click a resource, application, or general group and click View Details.
	3. Click Alerts in the General section.
	4. Double-click the alert then click Acknowledge in the alert details pane.
	5. To acknowledge multiple alerts, press Ctrl and click the alerts. Right-click the alerts and click View Alerts . Review the summary details of the alerts in the details pane, then click Acknowledge .
	To acknowledge all alerts, go to Home > Alerts , and then click Acknowledge all alerts in the Actions menu.
View all the alerts for resources, applications, and groups.	Go to Home > Alerts .
View all the alerts for a specific resource type.	 Go to the resource list page for the resource type that you want to view. For example, go to Network > Switches. Click the Alerts tab.
View the alerts for a specific resource.	 Go to the resource list page for the resource that you want to view. For example, go to Storage > Block Storage Systems. Right-click a resource and click View Details. Click Alerts in the General section.
View the alerts for a specific application.	 Go to Groups > Applications. Right-click an application and click View Details. Click Alerts in the General section.
View the alerts for a specific general group.	 Go to Groups > General Groups. Right-click a general group and click View Details. Click Alerts in the General section.

Table 39. Viewing and administering alerts (continued)	
Actions	Steps
View the details of an alert.	Double-click the alert. The alert details are displayed in a pane. Click another alert to display its details. Alternatively, click the close icon X.
Create an alert policy. <u>Learn</u> more	 Go to Settings > Alert Policies . To create a policy with default alert definitions, click Create Policy. To create a policy by copying an existing policy, select the policy to copy, then click Actions > Copy Policy.
Modify an alert policy. <u>Learn</u> more	Go to Settings > Alert Policies . To modify a policy, double-click the policy.

Viewing and administering alert definitions

View and administer the alert definitions and notification settings for alert policies, resources, applications, and general groups.

Viewing and administering alert definitions

Table 40. Viewing and administering alert definitions	
Actions	Steps
Create alert definitions in an alert policy. <u>Learn more</u>	 Go to Settings > Alert Policies. To create a policy with default alert definitions, click Create Policy. To create a policy by copying an existing policy, select the policy to copy, then click Actions > Copy Policy.
Modify alert definitions and notification settings in an alert policy. Learn more	 Go to Settings > Alert Policies. To modify a policy, double-click the policy. To edit the alert definitions, click Edit Alert Definitions on the Alert Definitions tab. To edit the notification settings, click Edit Policy Notifications. Restriction: You cannot modify the default alert policies.
Define and edit the alert definitions or notification settings for a resource that is not managed by an alert policy. Learn more about defining alerts for resources Learn more about modifying notification settings	 Go to the list page for the resource. Double-click the resource for which you want to define alerts. To edit the alert definitions, click Alert Definitions in the General section of the resource details page, then click Edit Alert Definitions. To edit the notification settings, click Edit Notifications.

Table 40. Viewing and administering alert definitions (continued)	
Actions	Steps
Define and edit the alert definitions for an application or general group. Learn more	 For applications, go to Groups > Applications. For general groups, go to Groups > General Groups. Right-click an application or general group and click View Alert Definitions.
Disable an alert definition for a resource, application, or general group.	 Go to the list page for the resource. Right-click the resource for which you want to disable an alert, then click View Alert Definitions. Click Edit Alert Definitions. Click the check mark to remove the check mark from the alert definition. Click Save Changes.

Alert policies

Use the **Alert Policies** page as a central location to view and manage all your alert policies.

To view the **Alerts Policies** page, click **Settings > Alert Policies**.

The **Alert Policies** page lists the default policies for your monitored resources and custom policies that you create. The page shows the following information for each alert policy:

- The types of resource that are allowed in the policy.
- The number of resources that are managed. To see which policies you are using, you can sort the list of policies by the number of resources.
- The number of alerts that are defined.
- Any email addresses that are specified to be notified about alerts instead of the default policy contacts.

You can manage the alert definitions and notification settings of the following types of resource using policies:

- Block storage systems
- File storage systems
- Object storage systems
- Servers
- Hypervisors
- Switches
- Fabrics

Here are some important points about alert policies:

- Alert policies manage one type of resource only.
- A resource can be managed by only one alert policy.
- It is not a requirement for resources to be managed by an alert policy. A resource can have its own alert definitions and notification settings, independent of an alert policy.
- If a resource is managed by a policy, the resource cannot have alert definitions and notification settings that are independent of the policy. The alert definitions and notification settings that apply to the resource come from the policy.
- When you add a resource to be monitored by IBM Spectrum Control, it is added to a default alert policy automatically.

- Default policies with alerts already configured are available. You can create copies of the default policies and assign resources to the new policies. Your alerts are configured with the default settings.
- If you add a resource to a policy, any existing alert definitions for the resource are replaced by the alert definitions in that policy. You can't restore the original alert definitions for a resource after you move it into the policy.

To store a copy of the original alert definitions for a resource before you move it into a policy, create a policy based on the alert configuration of the resource. Then, move the resource to another policy. If you want to reapply the original alert definitions later, you can add the resource to the policy with the original alert configuration.

The **Resources by Policy** page lists the resources and the policies that manage the resources.

Tips:

- If you have both Cisco and Brocade switches in your storage environment, you can manage the alert definitions for the different types of switch with different policies. Create the policies from the existing alert definitions in your switches. Alternatively, create copies of the default switch policy and edit the definitions to suit your requirements.
- Default alert policies are not provided for agentless servers.

[Video] How to create alert policies

Watch a short video about how create an alert policy and view alert notifications in IBM Storage Insights Pro.

Defining notification settings for alerts

You can define the alert notification settings to determine the actions that are taken when alert conditions are detected for a resource. The settings are applied to all of the alert definitions that are specified for the resource.

1. To define notification settings for alerts, choose one of the following options:


Option	Steps
Define notification settings for alerts in a policy	a. Go to Settings > Alert Policies. b. Double-click the policy. c. Click Edit Policy Notifications.
Define notification settings for alerts for a resource that is not managed by a policy	a. Go to the resource list page for the resource. For example, to modify notification settings for alerts for a block storage system, go to Storage > Block Storage Systems . To modify notification settings for a switch, go to Network > Switches .
	b. Right-click the resource for which you want to define alerts, then click View Alert Definitions. c. Click Edit Notifications in the Resource Notifications area.

2. If you want to send email notifications of alert violations, enter the email addresses in the **Email Addresses** field.

Tip: If you enter an email address in the **Email Addresses** field, only that email address receives notifications for the alert. The following contacts do not receive notifications:

- Any email addresses that are specified as policy contacts, if the alert is in an alert policy.
- Any global email addresses that are specified for alert notifications. To view the global alert notification addresses, go to **Settings** > **Notification Settings**.

3. To specify other notification actions, click the switch icon details required for the action.



You can specify that the following notification actions are taken when alert conditions are detected on monitored resources:

Run script

Run a script when an alert is triggered for the condition. Use a script to call external programs or run commands that take action as the result of an alert. By using a script, you can automatically address potential storage issues when they are detected to avoid unplanned downtime or performance bottlenecks. Learn more.

Restriction: You can only specify to run a script for an individual alert; you cannot specify to run a script for all alerts in an alert policy.

SNMP

Generate SNMP trap messages to any network management station (NMS), console, or terminal when an alert condition is detected. System administrators must set up their SNMP trap ringer with the provided management information base (MIB) files to receive SNMP traps from the product.

Netcool / OMNIbus

Send alert notifications to a Netcool server or OMNIbus EIF probe server within your environment that was configured to receive IBM Spectrum Control alerts.

Windows event log or UNIX syslog

Write alert messages to the OS log. If you already have an administrator monitoring OS logs, this method is a way to centralize your priority messages for quick notification and viewing.

4. Click Save Changes.

Alert severities

IBM Spectrum Control determines the severity of alert conditions that it detects on monitored resources. Use the severity level to help determine the priority in which you resolve alerts.

Each severity is represented by an icon in the web-based GUI.

Icon	Status	Description
8	Critical	An alert with a Critical severity represents a serious problem on a resource or on its internal resources. Resolve these problems as soon as possible. Review the condition that triggered an alert for more information about the problem.
16	Critical - Acknowledged	An alert with a Critical severity was acknowledged. A Critical - Acknowledged severity indicates that an alert was reviewed and is either resolved or can be ignored.
A	Warning	An alert with a Warning severity represents potential problems on a resource or on its internal resources. Resolve these problems after you fix any critical alerts. Review the condition that triggered an alert for more information about the problem.
1	Warning - Acknowledged	An alert with a Warning severity was acknowledged. A Warning - Acknowledged severity indicates that an alert was reviewed and is either resolved or can be ignored.
i	Informational	An alert with an Informational severity does not represent a problem, but is intended to provide information about actions related to a resource.

Icon	Status	Description
	Informational - Acknowledged	An alert with an Informational severity was acknowledged. An Informational - Acknowledged severity indicates that an alert was reviewed and can be ignored.

Defining alerts

Define how and when you are alerted to changes in your storage environment. You can define alerts for alert policies, for individual resources, or for sets of resources that are included in an application or general group.

Configure IBM Spectrum Control to send alert notifications by using email, SNMP traps, or a Tivoli Netcool/OMNIbus server.

Related reference

Alert notifications and actions

Specify how you are notified when alert conditions are detected on resources, and define actions to take as a result of those alerts. These settings are defined globally for all resources, and can be overridden for a specific alert definition, for all alert definitions that apply to a specific resource, or for an alert policy.

Defining alerts for resources

Define alerts for changes in the configuration, attributes, and performance of monitored resources.

Defining alert definitions for general attributes and capacity changes

You can define alerts that are triggered when the attributes or capacity of a resource changes. Attributes represent the key properties and configuration of a resource, such as status, versions, removals, discoveries, and data collection status. Capacity represents storage statistics such as available capacity, used capacity, drive capacity, reserved capacity, and more.

Asset, capacity, and configuration metadata is aggregated and collected when probes collect the metadata from the resources. By default, the metadata that is collected from storage systems is refreshed every 24 hours. Define alerts to track daily changes to the attributes of a storage system.

1. To define alerts for resources, choose one of the following options:

Option	Steps
Define alerts for a policy	 a. Go to Settings > Alert Policies. b. Double-click the policy. c. Click Edit Alert Definitions on the Alert Definitions tab.
Define alerts for a resource that is not managed by a policy	a. Go to the resource list page for the resource. For example, to define alerts for a block storage system, go to Storage > Block Storage Systems . To define alerts for a switch, go to Network > Switches .
	b. Right-click the resource for which you want to define alerts, then click View Alert Definitions.
	c. Click Edit Alert Definitions .

- 2. Click the type of resource that you want to alert on. For example, click **Storage System**.
- 3. Click the category of the attributes that you want to alert on.

Category	Description
General	Attributes for the key properties of a resource, such as status, version changes, removals, discoveries, state, and data collection status.
Capacity	Attributes for capacity statistics of a resource, such as available capacity, used capacity, drive capacity, reserved capacity, and more.

Restriction: Not all categories are available for some resources.

- 4. To enable the alert for an attribute, click the check mark for the attribute.

 The advanced options for the alert, such as notification frequency, are displayed.
- 5. Specify the conditions for generating an alert for an attribute.

Conditions can include operators such as *greater than or equal to*, or *less than or equal to*. Conditions can also include storage values and time values.

For example, for a capacity attribute such as Available Capacity, you can specify that an alert is generated when the amount of available capacity on a resource's pools is less than or equal to 50 GiB.



Tips:

- Not all attributes require conditions to generate an alert. For example, you can enable an alert for the Deleted Volume attribute, but you don't need to specify any conditions.
- Some attributes use operators such as *is*, *is not*, *contains*, and *changes*. For example, for the Firmware attribute for a DS8000 you can select the operator Contains and enter R5 in the value field. An alert is triggered if the firmware is at the R5 level rather than at a later version such as R6.1, R6.2, or R6.3. You can use this alert definition if you want to be notified when the firmware for a storage system is reverted to a previous version.
- 6. Assign a severity to an alert.

Assigning a severity can help you more quickly identify and address the critical conditions that are detected on resources. The severity that you assign depends on the guidelines and procedures within your organization. Default assignments are provided for each alert.

Option	Description
⊗ Critical	Alert is critical and needs to be resolved. For example, alerts that notify you when the amount of available capacity on a file system falls below a specified threshold.
/L Warning	Alerts that are not critical, but represent potential problems. For example, alerts that notify you when the status of a data collection job is not normal.
Informational	Alerts that might not require any action to resolve and are primarily for informational purposes. For example, alerts that are generated when a new pool is added to a storage system.

7. Optional: If you want to send email notifications of alert violations to contacts other than the policy contacts or global alert notification addresses, enter the email addresses in the **Email Override** field.

Tip: If you enter an email address in the **Email Override** field, only that email address receives notifications for the alert. The following contacts do not receive notifications:

- Any email addresses that are specified as policy contacts, if the alert is in an alert policy.
- Any global email addresses that are specified for alert notifications. To view the global alert notification addresses, go to **Settings** > **Notification Settings**.
- 8. Optional: Click **View Additional Options** to specify how frequently you are notified of alerts. Use these settings to avoid triggering too many alerts for some conditions.
- 9. Optional: Click **View Additional Options** to specify that the following actions are taken when alert conditions are detected on monitored resources:

Run script

Run a script when an alert is triggered for the condition. Use a script to call external programs or run commands that take action as the result of an alert. By using a script, you can automatically address potential storage issues when they are detected to avoid unplanned downtime or performance bottlenecks. Learn more.

Netcool / OMNIbus

Send alert notifications to a Netcool server or OMNIbus EIF probe server within your environment that was configured to receive IBM Spectrum Control alerts.

SNMP

Generate SNMP trap messages to any network management station (NMS), console, or terminal when an alert condition is detected. System administrators must set up their SNMP trap ringer with the provided management information base (MIB) files to receive SNMP traps from the product.

Windows event log or UNIX syslog

Write alert messages to the OS log. If you already have an administrator monitoring OS logs, this method is a way to centralize your priority messages for quick notification and viewing.

10. Optional: + Duplicate an alert.

Use this action when you want to define another alert for the same attribute but with different conditions and settings.

Duplicating alerts can be helpful in the following situations:

 When you want to generate separate warning alerts and critical alerts for different conditions on the same attribute.

For example, for a capacity attribute such as Available Capacity, you might want to define the following alerts:

- Define a warning alert 4 to be generated when the amount of available capacity on a resource's pools is less than or equal to 50 GiB.
- Duplicate the alert, but this time, specify a critical severity
 when the amount of available capacity on a resource's pools is less than or equal to 10 GiB.
- When you want to send, alert notifications to different people based on the severity of an alert.

In the previous example for the Available Pool Space attribute, you can configure the notification settings so that warning alerts are sent to junior administrators, while critical alerts are sent to more senior administrators.

11. Click Save Changes.

To view all the alerts generated by IBM Spectrum Control, go to **Home > Alerts** in the GUI.

Related concepts

Collecting data

Determining the data that you want to gather about storage resources is critical to helping you implement a storage management strategy. IBM Spectrum Control provides two different jobs for collecting data about resources.

Alerting

Specify conditions that trigger alerts and the actions to take when those alerts are triggered, such as notify an email address. Use alert policies to define those alert conditions and notification settings for a group of resources.

Related tasks

Viewing capacity alerts and violations

View the alerts that were triggered when the capacity of a resource changes and reaches a threshold. For example, you can view alerts that are generated when the measured value of a capacity metric for a pool meets the conditions for generating an alert.

Related reference

Alert notifications and actions

Specify how you are notified when alert conditions are detected on resources, and define actions to take as a result of those alerts. These settings are defined globally for all resources, and can be overridden for a specific alert definition, for all alert definitions that apply to a specific resource, or for an alert policy.

Defining alert definitions for performance changes

You can define alerts that are triggered when the performance of a resource falls outside a specified threshold.

1. To define alerts for resources, choose one of the following options:

Option	Steps
Define alerts for a policy	a. Go to Settings > Alert Policies. b. Double-click the policy. c. Click Edit Alert Definitions on the Alert Definitions tab.
Define alerts for a resource that is not managed by a policy	a. Go to the resource list page for the resource. For example, to define alerts for a block storage system, go to Storage > Block Storage Systems. To define alerts for a switch, go to Network > Switches.
	 b. Right-click the resource for which you want to define alerts, then click View Alert Definitions. c. Click Edit Alert Definitions.

- 2. Click the type of resource that you want to alert on. For example, click **Storage System**.
- 3. Click the **Performance** category.
- 4. To enable the alert for a performance metric, click the check mark for the metric. If the metric that you want is not displayed, click **Add Metrics** then select the metric you want.
- 5. Specify the conditions for generating an alert.

Conditions include an operator and a threshold value.

- a) Select an operator.
 - An operator determines if an alert is triggered when the performance of a resource is *greater than* or equal to or less than or equal to the specified threshold value.
- b) Enter a threshold value.
 - For example, to trigger an alert if the Total I/O Rate for a storage system is greater than or equal to 500 ops/s, enter the value 500.

Tips for threshold values:

- IBM Spectrum Control provides recommended values for threshold values that do not vary much between environments. For example, the default threshold values for Port Send Bandwidth Percentage are greater than or equal to 75% for warning alerts, and greater than or equal to 85% for critical alerts.
 - However, for metrics that measure throughput and response times, thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. In these cases, there are no recommended values. To help determine threshold values for a resource, collect performance data over time to establish a baseline of the normal and expected performance behavior for that resource. After you determine a set of baseline values, define alerts to trigger if the measured performance behavior falls outside the normally expected range.
- For some metrics, lower values might indicate more stress and higher values might indicate idle behavior. For example, a lower threshold value for the Cache Holding Time Threshold metric might indicate a performance problem.

6. Optional: Click **View Performance** to view a chart of the performance of the resource. Use the chart to evaluate the current and historical performance of a resource to help determine the threshold value for an alert.

In the performance chart, a yellow horizontal line represents the threshold value for warning alerts; a red horizontal line represents the threshold value for critical alerts.

To customize the chart, click **Top 10** or **Bottom 10** to show resources according to their performance, click a time period, and change the start and end dates for the data that is displayed.

7. Assign a severity to an alert.

Assigning a severity can help you more quickly identify and address the critical conditions that are detected on resources. The severity that you assign depends on the guidelines and procedures within your organization. Default assignments are provided for each alert.

Option	Description
⊗ Critical	Assign this severity to alerts that are critical and need to be resolved. For example, assign a critical severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 85%.
/\ Warning	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 75% but less than 85%.
i Informational	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes.

8. Optional: If you want to send email notifications of alert violations to contacts other than the policy contacts or global alert notification addresses, enter the email addresses in the **Email Override** field.

Tip: If you enter an email address in the **Email Override** field, only that email address receives notifications for the alert. The following contacts do not receive notifications:

- Any email addresses that are specified as policy contacts, if the alert is in an alert policy.
- Any global email addresses that are specified for alert notifications. To view the global alert notification addresses, go to **Settings** > **Notification Settings**.
- 9. Optional: Click **View Additional Options** to specify how frequently you are notified of alerts. Use these settings to avoid triggering too many alerts for some conditions.
- 10. Optional: Click **View Additional Options** to specify that the following actions are taken when alert conditions are detected on monitored resources:

Run script

Run a script when an alert is triggered for the condition. Use a script to call external programs or run commands that take action as the result of an alert. By using a script, you can automatically address potential storage issues when they are detected to avoid unplanned downtime or performance bottlenecks. Learn more.

Netcool / OMNIbus

Send alert notifications to a Netcool server or OMNIbus EIF probe server within your environment that was configured to receive IBM Spectrum Control alerts.

SNMP

Generate SNMP trap messages to any network management station (NMS), console, or terminal when an alert condition is detected. System administrators must set up their SNMP trap ringer with the provided management information base (MIB) files to receive SNMP traps from the product.

Windows event log or UNIX syslog

Write alert messages to the OS log. If you already have an administrator monitoring OS logs, this method is a way to centralize your priority messages for quick notification and viewing.

11. Optional: + Duplicate an alert.

Use this action when you want to define another alert for the same metric but with different conditions and settings.

Duplicating alerts can be helpful in the following situations:

• When you want to generate separate warning alerts and critical alerts for different thresholds on the same metric.

For example, for the CRC Error Rate metric for ports, you might want to define the following alerts:

- Define a warning alert to be generated when the number of frames per second that are received with cyclic redundancy check (CRC) errors is greater than or equal to .01 counts per second.
- Duplicate the alert, but this time, specify a critical severity when the CRC error rate is greater than or equal to .033 counts per second.
- · When you want to send alert notifications to different people based on the severity of an alert.

In the previous example for the CRC Error Rate metric, you can configure the notification settings so that warning alerts are sent to junior administrators, while critical alerts are sent to more senior administrators to resolve.

12. Click Save Changes.

To view all the alerts generated by IBM Spectrum Control, go to Home > Alerts in the GUI.

Tip: If a performance monitor is already collecting data about a resource when you add, modify, or remove a performance alert for that resource, changes are applied dynamically. You do not have to stop and restart the performance monitor to apply the changes. A confirmation message is recorded in the log of the performance monitor when the alert is updated.

Related reference

Alert notifications and actions

Specify how you are notified when alert conditions are detected on resources, and define actions to take as a result of those alerts. These settings are defined globally for all resources, and can be overridden for a specific alert definition, for all alert definitions that apply to a specific resource, or for an alert policy.

Defining custom alerts for resources

You can define alerts that are triggered when two or more changes occur in the attributes, capacity, and performance of resources.

To define a custom alert, select the general attributes, capacity, and performance metrics that you want to combine to trigger an alert and specify their conditions and threshold values. You can combine conditions for the resource and its internal resources into a custom alert. The alert is triggered when the conditions for the attributes and capacity of the resource are met, and the performance of the resource falls outside the threshold values.

For example, you can create a custom alert that notifies you when the overall response time for the volumes on a SAN Volume Controller system is worse than 20 milliseconds per operation *and* the system CPU utilization on the nodes on the system is greater than 70%. The Overall Response Time is a metric that measures the average number of milliseconds that it takes to service each I/O operation on a volume. The System CPU Utilization is a metric that measures the average percentage of time that the processors on nodes are busy doing system I/O tasks.

1. To define alerts for resources, choose one of the following options:

Option	Steps
Define alerts for a	a. Go to Settings > Alert Policies .
policy	b. Double-click the policy.
	c. Click Edit Alert Definitions on the Alert Definitions tab.

Option	Steps
Define alerts for a resource that is not managed by a policy	a. Go to the resource list page for the resource. For example, to define alerts for a block storage system, go to Storage > Block Storage Systems . To define alerts for a switch, go to Network > Switches .
	 b. Right-click the resource for which you want to define alerts, then click View Alert Definitions. c. Click Edit Alert Definitions.

- 2. Click Custom.
- 3. ① Click the create alert icon, then enter a name for the alert.
- 4. Assign a severity to the alert.

Assigning a severity can help you more quickly identify and address the critical conditions that are detected on resources. The severity that you assign depends on the guidelines and procedures within your organization.

Option	Description
⊗ Critical	Assign this severity to alerts that are critical and need to be resolved. For example, assign a critical severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 85%. The default severity for custom alerts is critical.
M Warning	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 75% but less than 85%.
Informational	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes.

- 5. Select a component, category, and group for the alert.
 For example, select Storage System, Capacity, and Available Capacity.
- 6. To generate an alert for a general or capacity attribute, specify the conditions for the alert.

Conditions can include operators such as *greater than or equal to*, or *less than or equal to*. Conditions can also include storage values and time values.

For example, for a capacity attribute such as Available Capacity, you can specify that an alert is generated when the amount of available capacity on a resource's pools is less than or equal to 50 GiB.



Tips:

- Not all attributes require conditions to generate an alert. For example, you can enable an alert for the Deleted Volume attribute, but you don't need to specify any conditions.
- Some attributes use operators such as *is, is not, contains*, and *changes*. For example, for the Firmware attribute for a DS8000 you can select the operator Contains and enter R5 in the value field. An alert is triggered if the firmware is at the R5 level rather than at a later version such as R6.1, R6.2, or R6.3. You can use this alert definition if you want to be notified when the firmware for a storage system is reverted to a previous version.
- 7. To generate an alert for a performance metric, specify the conditions for the alert.

Conditions include an operator and a threshold value.

a) Select an operator.

An operator determines if an alert is triggered when the performance of a resource is *greater than* or equal to or less than or equal to the specified threshold value.

b) Enter a threshold value.

For example, to trigger an alert if the Total I/O Rate for a storage system is greater than or equal to 500 ops/s, enter the value 500.

Tips for threshold values:

• IBM Spectrum Control provides recommended values for threshold values that do not vary much between environments. For example, the default threshold values for Port Send Bandwidth Percentage are greater than or equal to 75% for warning alerts, and greater than or equal to 85% for critical alerts.

However, for metrics that measure throughput and response times, thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. In these cases, there are no recommended values. To help determine threshold values for a resource, collect performance data over time to establish a baseline of the normal and expected performance behavior for that resource. After you determine a set of baseline values, define alerts to trigger if the measured performance behavior falls outside the normally expected range.

- For some metrics, lower values might indicate more stress and higher values might indicate idle behavior. For example, a lower threshold value for the Cache Holding Time Threshold metric might indicate a performance problem.
- 8. Optional: Click **View Performance** to view a chart of the performance of the resource. For performance attributes, use the chart to evaluate the current and historical performance of a resource to help determine the threshold value for an alert.

In the performance chart, a yellow horizontal line represents the threshold value for warning alerts; a red horizontal line represents the threshold value for critical alerts.

To customize the chart, click **Top 10** or **Bottom 10** to show resources according to their performance, click a time period, and change the start and end dates for the data that is displayed.

9. + Duplicate an alert.

Use this action to add a second and subsequent condition to the alert.

- 10. Repeat steps "5" on page 165 "9" on page 166to add more conditions to the alert.
- 11. Optional: If you want to send email notifications of alert violations to contacts other than the policy contacts or global alert notification addresses, enter the email addresses in the **Email Override** field.

Tip: If you enter an email address in the **Email Override** field, only that email address receives notifications for the alert. The following contacts do not receive notifications:

- Any email addresses that are specified as policy contacts, if the alert is in an alert policy.
- Any global email addresses that are specified for alert notifications. To view the global alert notification addresses, go to Settings > Notification Settings.
- 12. Optional: Click **View Additional Options** to specify how frequently you are notified of alerts. Use these settings to avoid triggering too many alerts for some conditions.
- 13. Optional: Click **View Additional Options** to specify that the following actions are taken when alert conditions are detected on monitored resources:

Run script

Run a script when an alert is triggered for the condition. Use a script to call external programs or run commands that take action as the result of an alert. By using a script, you can automatically address potential storage issues when they are detected to avoid unplanned downtime or performance bottlenecks. Learn more.

Netcool / OMNIbus

Send alert notifications to a Netcool server or OMNIbus EIF probe server within your environment that was configured to receive IBM Spectrum Control alerts.

SNMP

Generate SNMP trap messages to any network management station (NMS), console, or terminal when an alert condition is detected. System administrators must set up their SNMP trap ringer with the provided management information base (MIB) files to receive SNMP traps from the product.

Windows event log or UNIX syslog

Write alert messages to the OS log. If you already have an administrator monitoring OS logs, this method is a way to centralize your priority messages for quick notification and viewing.

14. Click Save Changes.

To view all the alerts generated by IBM Spectrum Control, go to Home > Alerts in the GUI.

Related reference

Alert notifications and actions

Specify how you are notified when alert conditions are detected on resources, and define actions to take as a result of those alerts. These settings are defined globally for all resources, and can be overridden for a specific alert definition, for all alert definitions that apply to a specific resource, or for an alert policy.

Scenarios for custom alerts

Use custom alerts to combine attributes and conditions from multiple resources into a single alert. By creating a custom alert, you can analyze multiple configuration, capacity, and performance conditions together to determine whether an urgent situation occurred in your storage.

The storage requirements of your environment determine the custom alerts that you create. For example, your storage systems might run critical production applications on tier 1 storage. In this case, you don't want the performance of the tier 1 storage to fall below a certain threshold. To be notified when that situation occurs, you can create a custom alert that checks if the overall response time of tier 1 storage is too high.

Tip: You can define alerts for storage systems, servers, hypervisors, fabrics, and switches and their internal resources. However, the resources in an alert must be along the same data path. For example, if you create an alert for a storage system volume, resources along the same data path include the pool that volume is part of, the switch ports that are used to access that volume, and the host that maps the volume.

Use the following example scenarios as guides to help you create custom alerts for your environment.

- "Receive alerts when the response time of storage on a specific tier is too high" on page 167
- "Receive alerts when the response time of volumes is too high during times of active I/O" on page 168
- "Receive alerts when the response time of volumes is too high, but do not generate these alerts when batch and backup jobs are running" on page 169
- "Receive alerts if a port is being used for both inter-node communication and host I/O exchanges" on page 170
- "Receive alerts for link resets that are not associated with link initialization" on page 171
- "Receive alerts for invalid word transmissions that are not associated with link initialization" on page 172

Receive alerts when the response time of storage on a specific tier is too high


Your storage systems run both critical production applications and noncritical test applications. The production applications use tier 1 storage, while the test applications use storage on tiers 2 and 3.

To ensure consistent, top performance for tier 1 storage, you want to be notified when its response time is higher than 6ms/op so that you can resolve the bottleneck. However, to avoid too many alerts, you do not want to receive notifications when the response time of tier 2 or 3 storage exceeds 6ms/op.

Solution

Define a custom alert that checks if the volumes used by an application are in Tier 1 pools and if their Overall Response Time is higher than 6 ms/op.

For a storage system, set up a custom alert with the following attributes and conditions:




Receive alerts when the response time of volumes is too high during times of active I/O

You care about high read response times on your volumes, but they can be caused by cache misses when there is only a trickle of I/O.

Solution

Define a custom alert with volume-level thresholds that combines checks for response times and I/O.



Receive alerts when the response time of volumes is too high, but do not generate these alerts when batch and backup jobs are running

To ensure the consistent, top performance of your volumes, you want to be notified when their response times are becoming too high. However, to avoid too many alerts, you do not want to receive notifications when batch and backup jobs are running on your storage. You understand that these jobs can cause an expected spike in response times and do not require action on your part.

Solution

Define a custom alert that checks if the Read Response Time of volumes exceeds an amount that is more than expected in your environment and the Read Transfer Size is less than 256 KiB/op. Typically, read transfer sizes greater than 256KiB/op indicate that batch or backup jobs are running in the background.

For a storage system, set up a custom alert with the following attributes and conditions:




Receive alerts if a port is being used for both inter-node communication and host I/O exchanges

You want to avoid potential bottlenecks by ensuring that storage system ports aren't being used for both inter-node communication in the local cluster and for I/O exchanges to host computers. You can also use this custom alert to check for adherence to best practices that are related to configuring ports for nodes with 8 or more ports. It does not apply to nodes that contain only 4 ports.

Solution

Define a custom alert that checks if the I/O rate for ports indicates exchanges between local nodes and hosts. For a storage system, set up a custom alert with the following attributes and conditions:



Attribute	Condition
Total Port-to-Host I/O Rate (Storage System > Ports > Performance)	Greater than or equal to (>=) .01 ops/s
Total Port-to-Local Node I/O Rate (Storage System > Ports > Performance)	Greater than or equal to (>=) .01 ops/s

Tip: Optionally, you can define other custom alerts to be notified of this situation, depending on your storage requirements. For example:


Attribute	Condition
Total Port-to-Disk I/O Rate (Storage System > Ports > Performance)	Greater than or equal to (>=) .01 ops/s
Total Port-to-Local Node I/O Rate (Storage System > Ports > Performance)	Greater than or equal to (>=) .01 ops/s
Total Port-to-Remote I/O Rate (Storage System > Ports > Performance)	Greater than or equal to (>=) .01 ops/s

Receive alerts for link resets that are not associated with link initialization

You want to identify link resets that are generated in response to hardware failures or link congestion. Link Resets generated by link initialization are ignored.

Solution

Define a custom alert that checks if link resets occur and if those resets are not associated with a link initialization. For a switch, set up a custom alert with the following attributes and conditions:



Attribute	Condition
Link Reset Received Rate (Storage System > Ports > Performance)	Greater than or equal to (>=) .01 cnt/s
OR	
Link Transmitted Received Rate (Storage System > Ports > Performance)	
Sync Loss (Storage System > Ports > Performance)	Less than or equal to (<=) 0 cnt/s
Signal Loss (Storage System > Ports > Performance)	Less than or equal to (<=) 0 cnt/s

Receive alerts for invalid word transmissions that are not associated with link initialization

You want to identify invalid transmission words that are generated because of poor link quality. Poor or marginal link quality can be caused by a bad SFP, HBA, or cable. Invalid transmission words that are generated by link initialization are ignored.

Solution

Define a custom alert that checks if invalid word transmissions occur and if those resets are not associated with a link initialization. For a switch, set up a custom alert with the following attributes and conditions:



Attribute	Condition
Invalid Transmission Word Rate (Storage System > Ports > Performance)	Greater than or equal to (>=) .01 cnt/s
Sync Loss (Storage System > Ports > Performance)	Less than or equal to (<=) 0 cnt/s
Signal Loss (Storage System > Ports > Performance)	Less than or equal to (<=) 0 cnt/s

Defining alerts for applications

Define alerts for changes in the configuration, attributes, and performance of the servers, volumes, filesets, and shares in your application.

You can quickly identify and address the critical conditions that are detected on the servers, volumes, filesets, and shares in your application by using alerts. Application alerts can be helpful in the following situations:

· When you want to apply different thresholds to internal resources of the same type on a storage system.

For example, you have production applications and test applications that use volumes on a SAN Volume Controller. The production applications require response times of 6 milliseconds or less while the test

applications can tolerate response times up to 30 milliseconds. You can use application alerts to set separate response time thresholds for volumes used by the different applications, depending on the needs of that application.

• When you want to quickly define alerts for multiple resources of the same type. You can define alerts once for the application and the alerts apply to all the resources of that type in the application.

For example, if your application uses multiple servers, you can define the server alerts once for the application and the alerts apply to all the servers. If you later add more servers to the application, the existing application alerts apply to those servers also.

Tips for working with alerts in application hierarchies:

- When you define an application alert for a resource such as a volume, the alert applies to all the resources of that type that belong to that application and all child applications.
- If you define an application alert for a resource such as a server, and the server also belongs to child applications, the alert is generated once at the parent application level. Separate alerts are not generated for each of the child applications that contain the server.
- If a child application has the same alert as a parent application but with different conditions, separate alerts are generated for the child application for the different alert conditions.

Related concepts

Applications

View information about storage area network resources that you can model by using applications in your business environment. The monitoring and management of applications enables you to perform capacity trending, health monitoring, and performance troubleshooting tasks.

Defining application alerts for attribute and capacity changes

You can define alerts that are triggered when the attributes or capacity of the servers, volumes, filesets, or shares in an application change.

Attributes represent the key properties and configuration of a resource, such as status, removals, discoveries, and data collection status. Capacity represents storage statistics such as available and used file system capacity, drive capacity, and volume used capacity.

To define an application alert for attribute and capacity changes, complete these steps:

- 1. In the menu bar, select **Groups** > **Applications**.
- 2. Right-click an application in the list and click View Alert Definitions.
- 3. Click Edit Alert Definitions.
- 4. Click the type of resource that you want to set an alert for. You can set an alert for one or more of the attributes of servers, volumes, filesets, or shares.
- 5. Click the category of the attributes that you want to set an alert for.

Category	Description
General	Attributes for the key properties of a resource, such as status, removals, discoveries, and data collection status.
Capacity	Attributes for capacity statistics of a resource, such as used capacity, available file system capacity, drive capacity, and volume used capacity.

- 6. To enable the alert for an attribute, click the check mark for the attribute.
- 7. Specify the conditions for generating an alert for an attribute.

Conditions can include operators such as *greater than or equal to*, or *less than or equal to*. Conditions can also include storage values and time values.

For example, for a capacity attribute such as Available Capacity, you can specify that an alert is generated when the amount of available capacity on a resource's pools is less than or equal to 50 GiB.



Tip: Not all attributes require conditions to generate an alert. For example, you can enable an alert for the New Fileset attribute for an IBM Spectrum Scale, but you don't need to specify any conditions.

8. Assign a severity to an alert to help you more quickly identify and address the critical conditions that are detected on resources.

The severity that you assign depends on the guidelines and procedures within your organization. Default assignments are provided for each alert.

Option	Description
⊗ Critical	Assign this severity to alerts that are critical and must be resolved. For example, assign a critical severity to alerts that notify you when the amount of available capacity on a file system falls below a specified threshold.
/L Warning	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the status of a data collection job is not normal.
Informational	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes. For example, assign an informational severity to alerts that are generated when a new fileset is added to a storage system.

9. Optional: If you want to send email notifications of alert violations to contacts other than the policy contacts or global alert notification addresses, enter the email addresses in the **Email Override** field.

Tip: If you enter an email address in the **Email Override** field, only that email address receives notifications for the alert. The following contacts do not receive notifications:

- Any email addresses that are specified as policy contacts, if the alert is in an alert policy.
- Any global email addresses that are specified for alert notifications. To view the global alert notification addresses, go to **Settings** > **Notification Settings**.
- 10. Optional: Click **View Additional Options** to specify how frequently you are notified of alerts. Use these settings to avoid triggering too many alerts for some conditions.
- 11. Optional: Click the icon + to duplicate an alert.

Use this action when you want to define another alert for the same attribute but with different criteria and settings.

Duplicating alerts can be helpful in the following situations:

• When you want to generate separate warning alerts and critical alerts for different conditions on the same attribute.

For example, for a capacity attribute such as Available Capacity, you might want to define the following alerts:

- Define a warning alert ¹ to be generated when the amount of available capacity on a resource's disks is less than or equal to 50 GiB.
- When you want to send alert notifications to different people based on the severity of an alert.

In the previous example for the Available Capacity attribute, you can configure the notification settings so that warning alerts are sent to junior administrators, while critical alerts are sent to more senior administrators to resolve.

12. Click Save Changes.

Related reference

Alert notifications and actions

Specify how you are notified when alert conditions are detected on resources, and define actions to take as a result of those alerts. These settings are defined globally for all resources, and can be overridden for a specific alert definition, for all alert definitions that apply to a specific resource, or for an alert policy.

Defining application alerts for performance metrics

You can define alerts that are triggered when the performance of the volumes that belong to an application fall outside a specified threshold.

To define an application alert for performance metrics, select the metric that you want to measure for the volumes in an application and specify a threshold value. When the performance of a volume falls outside the threshold, an alert is generated. For example, by using the Read Response Time metric you can define an alert that notifies you when the response time for a volume exceeds 20 milliseconds per read operation.

To define an application alert for performance metrics, complete these steps:

- 1. In the menu bar, select **Groups** > **Applications**.
- 2. Right-click an application in the list and click View Alert Definitions.
- 3. Click Edit Alert Definitions.
- 4. Click **Volumes** and then click **Performance**.
- 5. Click Add Metrics.
- 6. Select one or more metrics to alert on and click **OK**.
- 7. To enable the alert for a performance metric, click the check mark for the metric.
- 8. Specify the conditions for generating an alert.

Conditions include an operator and a threshold value.

- a) Select an operator.
 - An operator determines if an alert is triggered when the performance of a resource is *greater than* or equal to or less than or equal to the specified threshold value.
- b) Enter a threshold value.
 - For example, to trigger an alert if the Total I/O Rate for a storage system is greater than or equal to 500 ops/s, enter the value 500.

Tips for threshold values:

- IBM Spectrum Control provides recommended values for threshold values that do not vary much between environments. For example, the default threshold values for Port Send Bandwidth Percentage are greater than or equal to 75% for warning alerts, and greater than or equal to 85% for critical alerts.
 - However, for metrics that measure throughput and response times, thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. In these cases, there are no recommended values. To help determine threshold values for a resource, collect performance data over time to establish a baseline of the normal and expected performance behavior for that resource. After you determine a set of baseline values, define alerts to trigger if the measured performance behavior falls outside the normally expected range.
- For some metrics, lower values might indicate more stress and higher values might indicate idle behavior. For example, a lower threshold value for the Cache Holding Time Threshold metric might indicate a performance problem.
- 9. Assign a severity to an alert.

Assigning a severity can help you more quickly identify and address the critical performance conditions that are detected on volumes. The severity that you assign depends on the guidelines and procedures within your organization. Default assignments are provided for each alert.

Option	Description
⊗ Critical	Assign this severity to alerts that are critical and must be resolved. For example, you might assign a critical severity to alerts that notify you when the average Read Response Time for volumes is greater or equal to 20 ms/op.
/\ Warning	Assign this severity to alerts that are not critical, but represent potential problems. For example, you might assign a warning severity to alerts that notify you when the average Read Response Time for volumes is greater than 10 ms/op but less than 20 ms/op.
i Informational	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes.

10. Optional: Click **View Performance** to view a chart of the performance of the resource. Use the chart to evaluate the current and historical performance of a resource to help determine the threshold value for an alert.

In the performance chart, a yellow horizontal line represents the threshold value for warning alerts; a red horizontal line represents the threshold value for critical alerts.

To customize the chart, click **Top 10** or **Bottom 10** to show resources according to their performance, click a time period, and change the start and end dates for the data that is displayed.

11. Optional: If you want to send email notifications of alert violations to contacts other than the policy contacts or global alert notification addresses, enter the email addresses in the **Email Override** field.

Tip: If you enter an email address in the **Email Override** field, only that email address receives notifications for the alert. The following contacts do not receive notifications:

- Any email addresses that are specified as policy contacts, if the alert is in an alert policy.
- Any global email addresses that are specified for alert notifications. To view the global alert notification addresses, go to **Settings** > **Notification Settings**.
- 12. Optional: Click **View Additional Options** to specify how frequently you are notified of alerts. Use these settings to avoid triggering too many alerts for some conditions.
- 13. Optional: Click the icon + to duplicate an alert.

Use this action when you want to define another alert for the same metric but with different criteria and settings.

Duplicating alerts can be helpful in the following situations:

• When you want to generate separate warning alerts and critical alerts for different thresholds on the same metric.

For example, for the Read Response Time metric for volumes, you might want to define the following alerts:

- Define a warning alert 4 to be generated when the average time that it takes to service each read operation is greater than or equal to 10 milliseconds.
- When you want to send alert notifications to different people based on the severity of an alert.

In the previous example for the Read Response Time metric, you can configure the notification settings so that warning alerts are sent to junior administrators, while critical alerts are sent to more senior administrators to resolve.

14. Click Save Changes.

Related reference

Alert notifications and actions

Specify how you are notified when alert conditions are detected on resources, and define actions to take as a result of those alerts. These settings are defined globally for all resources, and can be overridden for a specific alert definition, for all alert definitions that apply to a specific resource, or for an alert policy.

Defining custom alerts for applications

Define a custom alert to combine multiple conditions from multiple resources for an application in a single alert. By creating a custom alert, you can analyze multiple configuration, capacity, and performance conditions together to determine whether an urgent situation occurred on the servers, volumes, filesets, and shares in an application.

For example, you can create a custom alert that notifies you when the used capacity on both volumes and filesets is greater than a specified amount.

Tip: In a custom alert, the resource types that you specify conditions for must be associated with the same type of top-level resource. For example, if you include conditions for storage system volumes, you can also include additional conditions for filesets and shares, but you cannot include conditions for servers.

- 1. In the menu bar, select **Groups** > **Applications**.
- 2. Right-click an application in the list and click View Alert Definitions.
- 3. Click Edit Alert Definitions.
- 4. Click Custom.
- 5. ① Click the create alert icon, then enter a name for the alert.
- 6. Assign a severity to the alert.

Assigning a severity can help you more quickly identify and address the critical conditions that are detected on resources. The severity that you assign depends on the guidelines and procedures within your organization.

Option	Description
⊗ Critical	Assign this severity to alerts that are critical and need to be resolved. For example, assign a critical severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 85%. The default severity for custom alerts is critical.
/\ Warning	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 75% but less than 85%.
ii Informational	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes.

- 7. Select a component, category, and group for the alert. For example, select Volumes, Capacity, and Used Capacity.
- 8. For general and capacity attributes, specify the criteria for generating an alert.

Use criteria such as greater than or equal to, less than or equal to, storage values, and time measurements to customize the conditions under which attributes generate alerts.

For example, for a capacity attribute such as Used Capacity, you can specify that an alert is generated when the amount of used capacity on a resource is more than or equal to 75%. The operator (less than or equal to) + a specified amount of space (75) + the unit of measurement (%) is the criteria that determines if an alert is generated.

Tips:

• Not all attributes require criteria for generating alerts. The category and type of an attribute determines whether you can specify criteria and the options that you can select.

- Some attributes can use the operators such as *is*, *is not*, and *changes*. For example, for the Used Inodes attribute for a fileset, select the operator **changes** to be notified if the current number of used inodes for a fileset changes.
- 9. For performance attributes, click **Add Metrics**. Specify the criteria for generating an alert. Criteria includes an operator and a threshold value.
 - a) Select an operator.

An operator determines if an alert is triggered when the performance of a resource is *greater than* or equal to or less than or equal to the specified threshold value.

b) Enter a threshold value.

For example, to trigger an alert if the Total I/O Rate for a volume is greater than or equal to 500 ops/s, enter the value 500.

Tips for threshold values:

- For metrics that measure throughput and response times, thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. In these cases, there are no recommended values. To help determine threshold values for a resource, collect performance data over time to establish a baseline of the normal and expected performance behavior for that resource. After you determine a set of baseline values, define alerts to trigger if the measured performance behavior falls outside the normally expected range.
- For some metrics, lower values might indicate more stress and higher values might indicate idle behavior. For example, a lower threshold value for the Cache Holding Time Threshold metric might indicate a performance problem.
- 10. Optional: Click **View Performance** to view a chart of the performance of the resource. Use the chart to evaluate the current and historical performance of a resource to help determine the threshold value for an alert.
 - In the performance chart, a yellow horizontal line represents the threshold value for warning alerts; a red horizontal line represents the threshold value for critical alerts.
 - To customize the chart, click **Top 10** or **Bottom 10** to show resources according to their performance, click a time period, and change the start and end dates for the data that is displayed.
- 11. Optional: If you want to send email notifications of alert violations to contacts other than the policy contacts or global alert notification addresses, enter the email addresses in the **Email Override** field.

Tip: If you enter an email address in the **Email Override** field, only that email address receives notifications for the alert. The following contacts do not receive notifications:

- Any email addresses that are specified as policy contacts, if the alert is in an alert policy.
- Any global email addresses that are specified for alert notifications. To view the global alert notification addresses, go to **Settings** > **Notification Settings**.
- 12. Optional: Click **View Additional Options** to specify how frequently you are notified of alerts. Use these settings to avoid triggering too many alerts for some conditions.
- 13. + Add another condition to the custom alert and repeat steps 7 11.

 If you do not want to add another condition, continue to the next step.
- 14. Click Save Changes.

Related reference

Alert notifications and actions

Specify how you are notified when alert conditions are detected on resources, and define actions to take as a result of those alerts. These settings are defined globally for all resources, and can be overridden for a specific alert definition, for all alert definitions that apply to a specific resource, or for an alert policy.

Defining alerts for general groups

Define alerts for changes in the configuration, attributes, and performance of the resources in your general groups.

Before you can define general group alerts for a resource type, you must first add resources of that type to the group or one of its subgroups. For example, if you want to set an alert for the attributes of a switch, you must first add one or more switches to your general group or one of its subgroups.

Tips for working with alerts in general group hierarchies:

- When you edit the alert definitions for a general group, you can edit the alert definitions for its subgroups at the same time. The resources for all subgroups of the general group are shown when you edit the alert definitions for a parent general group.
- If you define a general group alert for a resource such as a server, and the server also belongs to subgroups, the alert is generated once at the parent group level. Separate alerts are not generated for each of the subgroups that contain the server.
- If a subgroup has the same alert as a parent group but with different conditions, separate alerts are generated for the subgroup for the different alert conditions.

Related reference

General groups

View and administer the general groups that are used to alert on logically related storage resources. You can define alerts for a general group to notify you about changes in the configuration, attributes, and performance of the resources in the group.

Defining general group alerts for attribute and capacity changes

You can define alerts that are triggered when the attributes or capacity of the resources that belong to a general group change.

Attributes represent the key properties and configuration of a resource, such as status, removals, discoveries, and data collection status. Capacity represents storage statistics such as available and used file system capacity, drive capacity, and volume used capacity.

To define a general group alert for attribute and capacity changes, complete these steps:

- 1. In the menu bar, select **Groups** > **General Groups**.
- 2. Right-click a general group in the list and click **View Alert Definitions**.
- 3. Click Edit Alert Definitions.
- 4. Click the type of resource that you want to set an alert for.

You can set an alert for one or more of the attributes of the resource itself or its internal resources. For example, for a switch, you can set an alert for the attributes of the switch itself and for the attributes of the blades, inter-switch connections, and ports.

5. Click the category of the attributes that you want to alert for.

Category	Description
General	Attributes for the key properties of a resource, such as status, removals, discoveries, and data collection status.
Capacity	Attributes for capacity statistics of a resource, such as available and used capacity, drive capacity, file systems, volumes, and pools, and reserved capacity.

- 6. To enable the alert for an attribute, click the check mark for the attribute.
- 7. Specify the conditions for generating an alert for an attribute.

Conditions can include operators such as *greater than or equal to*, or *less than or equal to*. Conditions can also include storage values and time values.

For example, for a capacity attribute such as Available Capacity, you can specify that an alert is generated when the amount of available capacity on a resource's pools is less than or equal to 50 GiB.



Tips:

- Not all attributes require criteria for generating alerts. The category and type of an attribute determines whether you can specify criteria and the options that you can select.
- Some attributes can use the operators such as *is, is not,* and *contains*. For example, for the Firmware attribute on a DS8000 storage system, select the operator **contains**. Then, in the value field, type R5 to be notified if the firmware is at the R5 level rather than at a later version such as R6.3, R6.2, or R6.3.
- 8. Assign a severity to an alert to help you more quickly identify and address the critical conditions that are detected on resources.

The severity that you assign depends on the guidelines and procedures within your organization. Default assignments are provided for each alert.

Option	Description
⊗ Critical	Assign this severity to alerts that are critical and must be resolved. For example, assign a critical severity to alerts that notify you when the amount of available capacity on a file system falls below a specified threshold.
/\ Warning	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the status of a data collection job is not normal.
Informational	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes. For example, assign an informational severity to alerts that are generated when a new fileset is added to a storage system.

9. Optional: If you want to send email notifications of alert violations to contacts other than the policy contacts or global alert notification addresses, enter the email addresses in the **Email Override** field.

Tip: If you enter an email address in the **Email Override** field, only that email address receives notifications for the alert. The following contacts do not receive notifications:

- Any email addresses that are specified as policy contacts, if the alert is in an alert policy.
- Any global email addresses that are specified for alert notifications. To view the global alert notification addresses, go to **Settings** > **Notification Settings**.
- 10. Optional: Click **View Additional Options** to specify how frequently you are notified of alerts. Use these settings to avoid triggering too many alerts for some conditions.
- 11. Optional: Click the icon + to duplicate an alert.

Use this action when you want to define another alert for the same attribute but with different criteria and settings.

Duplicating alerts can be helpful in the following situations:

• When you want to generate separate warning alerts and critical alerts for different conditions on the same attribute.

For example, for a capacity attribute such as Available Capacity, you might want to define the following alerts:

- Define a warning alert 4 to be generated when the amount of available capacity on a resource's disks is less than or equal to 50 GiB.
- Duplicate the alert, but this time, specify a critical severity [™] when the amount of available capacity on a resource's disks is less than or equal to 10 GiB.
- · When you want to send alert notifications to different people based on the severity of an alert.

In the previous example for the Available Capacity attribute, you can configure the notification settings so that warning alerts are sent to junior administrators, while critical alerts are sent to more senior administrators to resolve.

12. Click Save Changes.

Related reference

Alert notifications and actions

Specify how you are notified when alert conditions are detected on resources, and define actions to take as a result of those alerts. These settings are defined globally for all resources, and can be overridden for a specific alert definition, for all alert definitions that apply to a specific resource, or for an alert policy.

Defining general group alerts for performance metrics

You can define alerts that are triggered when the performance of the storage systems or switches that belong to a general group fall outside a specified threshold.

To define a general group alert for performance metrics, select the metric that you want to measure and specify a threshold value. When the performance of that resource falls outside the threshold, an alert is generated. For example, you can define an alert that notifies you when the back-end response times for managed disks on a SAN Volume Controller exceed 35 milliseconds per read operation. The Overall Back-end Response Time is a metric that measures the average number of milliseconds that it takes to service each read operation on a managed disk.

To define a general group alert for performance metrics, complete these steps:

- 1. In the menu bar, select **Groups** > **General Groups**.
- 2. Right-click a general group in the list and click **View Alert Definitions**.
- 3. Click Edit Alert Definitions.
- 4. Click the type of resource that you want to alert on.
- 5. Click Performance and then click Add Metrics.
- 6. Select one or more metrics to alert on and click **OK**.
- 7. To enable the alert for a performance metric, click the check mark for the metric.
- 8. Specify the conditions for generating an alert.

Conditions include an operator and a threshold value.

a) Select an operator.

An operator determines if an alert is triggered when the performance of a resource is *greater than* or equal to or less than or equal to the specified threshold value.

b) Enter a threshold value.

For example, to trigger an alert if the Total I/O Rate for a storage system is greater than or equal to 500 ops/s, enter the value 500.

Tips for threshold values:

• IBM Spectrum Control provides recommended values for threshold values that do not vary much between environments. For example, the default threshold values for Port Send Bandwidth Percentage are greater than or equal to 75% for warning alerts, and greater than or equal to 85% for critical alerts.

However, for metrics that measure throughput and response times, thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. In these cases, there are no recommended values. To help determine threshold values for a resource, collect performance data over time to establish a baseline of the normal and expected performance behavior for that resource. After you determine a set of baseline values, define alerts to trigger if the measured performance behavior falls outside the normally expected range.

- For some metrics, lower values might indicate more stress and higher values might indicate idle behavior. For example, a lower threshold value for the Cache Holding Time Threshold metric might indicate a performance problem.
- 9. Optional: Click **View Performance** to view a chart of the performance of the resource. Use the chart to evaluate the current and historical performance of a resource to help determine the threshold value for an alert.

In the performance chart, a yellow horizontal line represents the threshold value for warning alerts; a red horizontal line represents the threshold value for critical alerts.

To customize the chart, click **Top 10** or **Bottom 10** to show resources according to their performance, click a time period, and change the start and end dates for the data that is displayed.

10. Assign a severity to an alert.

Assigning a severity can help you more quickly identify and address the critical performance conditions that are detected on resources. The severity that you assign depends on the guidelines and procedures within your organization. Default assignments are provided for each alert.

Option	Description
⊗ Critical	Assign this severity to alerts that are critical and must be resolved. For example, assign a critical severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 85%.
M Warning	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 75% but less than 85%.
Informational	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes.

11. Optional: If you want to send email notifications of alert violations to contacts other than the policy contacts or global alert notification addresses, enter the email addresses in the **Email Override** field.

Tip: If you enter an email address in the **Email Override** field, only that email address receives notifications for the alert. The following contacts do not receive notifications:

- Any email addresses that are specified as policy contacts, if the alert is in an alert policy.
- Any global email addresses that are specified for alert notifications. To view the global alert notification addresses, go to **Settings** > **Notification Settings**.
- 12. Optional: Click **View Additional Options** to specify how frequently you are notified of alerts. Use these settings to avoid triggering too many alerts for some conditions.
- 13. Optional: Click the icon + to duplicate an alert.

Use this action when you want to define another alert for the same metric but with different criteria and settings.

Duplicating alerts can be helpful in the following situations:

 When you want to generate separate warning alerts and critical alerts for different thresholds on the same metric.

For example, for the CRC Error Rate metric for ports, you might want to define the following alerts:

- Define a warning alert to be generated when the number of frames per second that are received with cyclic redundancy check (CRC) errors is greater than or equal to 0.01 counts per second.
- Duplicate the alert, but this time, specify a critical severity when the CRC error rate is greater than or equal to 0.033 counts per second.
- When you want to send alert notifications to different people based on the severity of an alert.

In the previous example for the CRC Error Rate metric, you can configure the notification settings so that warning alerts are sent to junior administrators, while critical alerts are sent to more senior administrators to resolve.

14. Click Save Changes.

Related reference

Alert notifications and actions

Specify how you are notified when alert conditions are detected on resources, and define actions to take as a result of those alerts. These settings are defined globally for all resources, and can be overridden for a specific alert definition, for all alert definitions that apply to a specific resource, or for an alert policy.

Defining custom alerts for general groups

Define a custom alert to combine multiple conditions from multiple resources for a general group in a single alert. By creating a custom alert, you can analyze multiple configuration, capacity, and performance conditions together to determine whether an urgent situation occurred on the resources in a group.

For example, you can create a custom alert that notifies you when the response times for volumes and pools in a group exceed a certain threshold and the performance monitors for storage systems are running without problems.

Tip: In a custom alert, the resource types that you specify conditions for must be associated with the same type of top-level resource. For example, if you include conditions for storage systems, you can also include additional conditions for pools and volumes, but you cannot include conditions for servers.

- 1. In the menu bar, select **Groups** > **General Groups**.
- 2. Right-click a general group in the list and click **View Alert Definitions**.
- 3. Click Edit Alert Definitions.
- 4. Click Custom.
- 5. Click Create Alert, then enter a name for the custom alert.
- 6. Assign a severity to the alert.

Assigning a severity can help you more quickly identify and address the critical conditions that are detected on resources. The severity that you assign depends on the guidelines and procedures within your organization.

Option	Description
⊗ Critical	Assign this severity to alerts that are critical and need to be resolved. For example, assign a critical severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 85%. The default severity for custom alerts is critical.
/\ Warning	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 75% but less than 85%.
i Informational	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes.

7. Select a component, category, and group for the alert. For example, select Volumes, Capacity, and Available Capacity.

8. For general and capacity attributes, specify the criteria for generating an alert.

Use criteria such as greater than or equal to, less than or equal to, storage values, and time measurements to customize the conditions under which attributes generate alerts.

For example, for a capacity attribute such as Used Capacity, you can specify that an alert is generated when the amount of used capacity on a resource is more than or equal to 75%. The operator (less than or equal to) + a specified amount of space (75) + the unit of measurement (%) is the criteria that determines if an alert is generated.

Tips:

- Not all attributes require criteria for generating alerts. The category and type of an attribute determines whether you can specify criteria and the options that you can select.
- Some attributes can use the operators such as *is, is not*, and *changes*. For example, for the Firmware attribute on a DS8000 storage system, select the operator **contains**. Then, in the value field, type R5 to be notified if the firmware is at the R5 level rather than at a later version such as R6.3, R6.2, or R6.3. This alert definition might be useful if you want to be notified when the firmware for a storage system was reverted to a previous version for some reason.
- 9. For performance attributes, specify the criteria for generating an alert.

Criteria includes an operator and a threshold value.

a) Select an operator.

An operator determines if an alert is triggered when the performance of a resource is *greater than* or equal to or less than or equal to the specified threshold value.

b) Enter a threshold value.

For example, to trigger an alert if the Total I/O Rate for a volume is greater than or equal to 500 ops/s, enter the value 500.

Tips for threshold values:

- For metrics that measure throughput and response times, thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. In these cases, there are no recommended values. To help determine threshold values for a resource, collect performance data over time to establish a baseline of the normal and expected performance behavior for that resource. After you determine a set of baseline values, define alerts to trigger if the measured performance behavior falls outside the normally expected range.
- For some metrics, lower values might indicate more stress and higher values might indicate idle behavior. For example, a lower threshold value for the Cache Holding Time Threshold metric might indicate a performance problem.
- 10. Optional: Click **View Performance** to view a chart of the performance of the resource. Use the chart to evaluate the current and historical performance of a resource to help determine the threshold value for an alert.
 - In the performance chart, a yellow horizontal line represents the threshold value for warning alerts; a red horizontal line represents the threshold value for critical alerts.
 - To customize the chart, click **Top 10** or **Bottom 10** to show resources according to their performance, click a time period, and change the start and end dates for the data that is displayed.
- 11. Optional: If you want to send email notifications of alert violations to contacts other than the policy contacts or global alert notification addresses, enter the email addresses in the **Email Override** field.

Tip: If you enter an email address in the **Email Override** field, only that email address receives notifications for the alert. The following contacts do not receive notifications:

- Any email addresses that are specified as policy contacts, if the alert is in an alert policy.
- Any global email addresses that are specified for alert notifications. To view the global alert notification addresses, go to Settings > Notification Settings.
- 12. +Add another condition to the custom alert and repeat steps 7 11.

If you do not want to add another condition, continue to the next step.

13. Optional: Click View Additional Options to specify how frequently you are notified of alerts.

Use these settings to avoid triggering too many alerts for some conditions.

14. Click Save Changes.

Related reference

Alert notifications and actions

Specify how you are notified when alert conditions are detected on resources, and define actions to take as a result of those alerts. These settings are defined globally for all resources, and can be overridden for a specific alert definition, for all alert definitions that apply to a specific resource, or for an alert policy.

Configuring alert notifications

Alerts can define notification actions that send email, generate Simple Network Management Protocol (SNMP) traps, or generate IBM Tivoli Netcool/OMNIbus events. To enable these notification actions, you must configure IBM Spectrum Control for email, SNMP, or Tivoli Netcool/OMNIbus alert notifications.

Tip: For alert notifications to generate SNMP traps or IBM Tivoli Netcool/OMNIbus events, first configure IBM Spectrum Control for these notification actions. Then, enable the actions as alert notifications for an alert policy, a resource, or an individual alert.

Configuring email alert notifications

To send alert notifications and reports, you must configure the email server.

To add or modify settings for the email server, you must be assigned the Administrator role.

The email server that is used to send alert notifications is also used to send reports. If you didn't configure the email server for alert notifications, you must configure the email server before you create reports.

IBM Spectrum Control ensures, when it sends alert notifications and reports by email to your email server, that it complies with the security standards that you configure on your email server. For example, if your email server requires authentication and TLS to send emails, then IBM Spectrum Control uses the authentication credentials that you provide when you set up the email server, and uses TLS to establish the connection with your email server. Because the email server that you set up is automatically trusted, you don't have any additional management tasks such as importing the server certificate into the keystores that are used by IBM Spectrum Control.

- 1. Click Settings > Notification Settings.
- 2. Type the host name, or IPv4, or IPv6 address of the email server.
- 3. Type the port number that is used by the email server to send alert notifications and reports.
- 4. Type the name of the user and the password that you want to use to authenticate with the email server.
- 5. Optional: Type the email address, such as the email address of the administrator, that you want to use for receiving replies to alert notifications and reports.

If you don't enter a value, and you entered an email address in the **User Name** field, then replies are sent to that address. Otherwise, replies are sent to no-reply@<Spectrum_Control_server_host_name>.

Tip: Before you save the settings for the email server, test the connection.

6. Optional: To specify the email addresses that you want to notify when an alert notification is generated, click **Global email notification settings**, and add each email address separated by a comma.

You can override the global setting for alert notifications by editing the notification settings for a specific alert policy, for a specific resource, or for a specific alert definition that is associated with a resource.

Restriction: Applies only to alert notifications.

7. Click Save.

Configuring SNMP alert notifications

You can define an alert to generate SNMP traps when an alert condition is detected on a monitored resource. To enable SNMP alert notifications, configure IBM Spectrum Control to direct the traps to at least one SNMP destination.

To modify alert notification settings, you must be assigned the Administrator role.

You must also be running an SNMP management application in your system environment. A Management Information Base (MIB) file is provided in the IBM Spectrum Control installation. You must set up the SNMP trap receiver with the MIB file to receive SNMP traps.

The MIB is automatically deployed when you install IBM Spectrum Control, and can also be found on the installation media.

Table 41. Location of MIB file	
Location of MIB file	Directory
Installation media	\data\snmp\tivoliSRM.mib
After IBM Spectrum Control is installed	<pre>installation_dir\data\snmp \tivoliSRM.mib</pre>

To configure SNMP alert notifications, complete the following steps:

- 1. Go to Settings > Notification Settings.
- 2. Click the **SNMP** tab.
- 3. Click **Edit**, and specify new settings or remove existing settings.
 - To specify new settings, complete the following steps:
 - a. Specify the following SNMP alert notification settings for each SNMP destination. You can specify up to two SNMP destinations.

Community

Specify the name of the SNMP community for sending SNMP traps. By default, the community is *public*.

IP address

Specify the host that is configured to receive SNMP traps. You can specify a host name, an IPv4 address, or an IPv6 address depending on what is supported within your environment.

Port

Specify the port number for receiving SNMP traps. SNMP trap messages are sent to this port when an alert condition is detected. By default, the port is set to 162.

- b. Click Save.
- To delete the current configuration settings without specifying new settings, click **Remove**.

Configuring Tivoli Netcool/OMNIbus alert notifications

You can define an alert to generate Tivoli Netcool/OMNIbus events when an alert condition is detected on a monitored resource. To enable Tivoli Netcool/OMNIbus notifications, configure IBM Spectrum Control to direct the events to a Tivoli Netcool/OMNIbus server.

To modify alert notification settings, you must be assigned the Administrator role.

In your system environment, you must be running a Tivoli Netcool/OMNIbus server that is configured to receive IBM Spectrum Control alerts.

To configure Tivoli Netcool/OMNIbus alert notifications, complete the following steps:

- 1. Go to Settings > Notification Settings.
- 2. Click **Netcool/OMNIibus** in the left-side navigation.
- 3. Click **Edit** and specify new settings or remove existing settings.
 - To specify new settings, complete the following steps:

a. Specify the following Tivoli Netcool/OMNIbus alert notification settings:

IP address

Specify the NetCool/OMNIbus server that is configured to receive notifications. You can specify a host name, and IPv4 address, or an IPv6 address. You cannot use a server that is configured for IPv6 only. The server must be IPv4 or dual stack enabled.

Port

Specify the port for receiving alert notifications. Alert notifications are sent to this port when an alert condition is detected. By default, the port is set to 5529.

- b. Click Save.
- To delete the current configuration settings without specifying new settings, click **Remove**.
- 4. Copy the EIF rules files that are provided with IBM Spectrum Control to the appropriate directory in the Tivoli Netcool/OMNIbus installation.

Rules file names		Copy rules files to this location in Tivoli Netcool/OMNIbus
tivoli_eif_tpc.rulestivoli_eif_tpc_tbsm.rules	extracted_installation_i mage/data/omnibus	<i>\$OMNIHOME</i> /probes/arch/

For more information about rules file, see the Tivoli Netcool/OMNIbus Knowledge Center.

Triggering conditions for alerts

Define alerts so that IBM Spectrum Control automatically notifies you when certain conditions or events are detected on monitored resources. Such conditions are the *triggering conditions* for the alert. The specific conditions that can trigger alerts depend on the type of resource that is being monitored.

Related concepts

Alerting

Specify conditions that trigger alerts and the actions to take when those alerts are triggered, such as notify an email address. Use alert policies to define those alert conditions and notification settings for a group of resources.

Related tasks

Adding resources

You must add resources to IBM Spectrum Control before you can collect data, generate reports, and manage storage that is related to those resources.

Triggering conditions for storage system alerts

You can set up IBM Spectrum Control so that it examines the attributes, capacity, and performance of a storage system and notifies you when changes or violations are detected.

Alerts can notify you of general changes and performance issues on the following resources:

- Storage Systems (performance)
- "Storage Systems (general changes)" on page 189
- "Triggering conditions for storage system internal resource alerts" on page 194

Important: Not all the attributes upon which you can alert are listed here. To view a complete list of attributes upon which you can alert, go to **Settings** > **Alert Policies**. Double-click a default policy for a storage system. Click **Edit Alert Definitions** on the **Alert Definitions** tab. View the attributes that are available in the general, capacity, and performance categories. Note that the attributes that are automatically configured for alerts in the default alert policies, or default alerts, have a status of Active.

In the tables, default alerts are marked with an asterisk (*).

Tips:

- The type of storage system determines which attributes and performance conditions are available for alerts. For example, triggering conditions for shares are available only for storage systems that are configured for file storage, such as Storwize V7000 Unified.
- For capacity attributes, you can generate alerts when the amount of storage is greater than, less than, or equal to a specified value. You can also determine the unit of measurement for the attribute, such as KiB, MiB, GiB, or TiB.
- If you are doing tasks where many volumes are being deleted, you might want to temporarily disable alerts that use the Deleted Volume attribute. For example, you might want to disable Deleted Volume alerts temporarily if you are doing maintenance tasks or decommissioning storage.

Storage Systems (performance)

Define alerts that notify you when the performance of a storage system falls outside a specified threshold. In alerts, you can specify conditions based on metrics that measure the performance of volumes, disks, ports, and nodes. By creating alerts with performance conditions, you can be informed about potential bottlenecks in your storage infrastructure.

Examples:

- You can define an alert to be notified when the average number of I/O operations per second for read and write operations on a storage system's volumes is greater than or equal to a specified threshold. Use this alert to be notified when the workload of a volume is high and you might need to balance that load across other volumes to improve overall performance.
- You can define an alert to be notified when the percentage of the average response time that can be attributed to delays from host systems is greater than or equal to a specified threshold. Use this alert to be notified of slow hosts and fabrics that might not be working efficiently.
- You can also define an alert that notifies you when a metric is less than a specified threshold, such as if you want to identify volumes that might be under used.

Tips:

- The type of storage system determines the metrics that can be alerted upon. For a list of the metrics that are available for each type of storage system, see "Performance metrics" on page 277.
- A performance monitor must collect data about a resource before IBM Spectrum Control can determine whether a threshold is violated and an alert is generated for a performance condition.

Best practice: When you set thresholds for performance conditions, try to determine the best value so you can derive the maximum benefit without generating too many false alerts. Because suitable thresholds are highly dependent on the type of workload that is being run, hardware configuration, the number of physical disks, exact model numbers, and other factors, there are no easy or standard default rules.

A recommended approach is to monitor the performance of resources for a number of weeks and by using this historical data, determine reasonable threshold values for each performance condition. After that is done, you can fine-tune the condition settings to minimize the number of false alerts.

Storage Systems (general changes)

Asset, capacity, and configuration metadata is aggregated and collected when probes collect storage system metadata. By default, metadata is collected once every 24 hours.

Table 42. Triggering attributes and conditions for general changes on storage systems		
General Attributes	Triggering Conditions for Attributes	
Firmware	The firmware version of the microcode on a storage system. For the DS-series of storage systems, this value represents the SEA version of the firmware.	
	To view information about the code bundles for the firmware versions of the DS-series, go to the <u>IBM support site</u> at http://www.ibm.com/support/entry/portal/Overview and search for code bundle information. An internet connection is required to access the support site.	
Last Successful Probe Last Successful Monitor	A specified amount of time has passed since a probe or performance monitor was able to collect data about a storage system. You can use this alert to be notified when up-to-date configuration, status, or performance data is not being collected about a storage system and its existing data might be stale. This situation might occur if the resource, network, or IBM Spectrum Control server is unavailable.	
Performance Monitor Status*	One of the following statuses is detected for a performance monitor:	
	Not Normal An error or warning occurred during data collection. This status indicates that a performance monitor did not collect any data, or only collected a partial set of data about a resource.	
	Warning A performance monitor completed, but did not collect a complete set of performance data. This status might occur if the resource was rebooted during data collection, no valid performance data was provided by the resource, or a communication error occurred with the resource or its associated agent.	
	Error A performance did not complete when it attempted to collect performance data about the resource. This status might occur if the resource cannot be reached during data collection, or if no configuration data is available for the resource.	
	For details about why a specific status occurred, check the log for the performance monitor. To check the log, go to the details page for a resource, click Data Collection , and select Actions > Open Logs in the Performance Monitor section on the Data Collection page.	

Table 42. Triggering attribute	able 42. Triggering attributes and conditions for general changes on storage systems (continued)	
General Attributes Triggering Conditions for Attributes		
Probe Status*	One of the following statuses is detected for a probe:	
	Not Successful An error or warning occurred during data collection. This status indicates that a probe did not collect any data, or only collected a partial set of data about a resource.	
	Warning A probe completed, but might not have collected a complete set of data. This status might occur if data cannot be collected about one or more of the internal resources of a resource.	
	Error (default) A probe did not complete when it attempted to collect asset data about the resource. This status might occur if the resource cannot be reached during data collection.	
	For details about why a specific status occurred, check the log for the probe. To check the log, go to the details page for a resource, click Data Collection , and select Actions > Open Logs in the Probe section on the Data Collection page.	
Status	One of the following statuses is detected for a storage system:	
	Not Normal An error or warning status was detected for a storage system.	
	Warning A warning status was detected for a storage system. This status might occur if a storage system comes online or if its version changes.	
	Error An error status was detected for a storage system. This status might occur if the cooling fans in a storage system are stopped and the internal temperature is too high or if a storage system goes offline.	

Table 43. Triggering attributes and conditions for capacity changes on storage systems		
Capacity Attributes	Triggering Conditions for Attributes	
Adjusted Used Capacity	The amount of capacity that can be used without exceeding the capacity limit. For example, you set a capacity limit of 80% for your storage systems. You want to get an informational alert when the adjusted used capacity exceeds 60% and a critical alert when the adjusted used capacity exceeds 80%. So, you define an informational alert with these parameters:	
	Adjusted Used Capacity ≥ 60%	
	And, you define a critical alert with these parameters:	
	Adjusted Used Capacity ≥ 80%	
Available Capacity (Previously known as Available Pool	The total amount of the space in the pools that is not allocated to the volumes in the pools. To calculate available capacity, the following formula is used:	
Space)	(pool capacity - used capacity)	
	For XIV systems, pool capacity is the physical capacity of the pools and does not include the provisioned capacity of the pools.	

Table 43. Triggering attributes and conditions for capacity changes on storage systems (continued)		
Capacity Attributes	Triggering Conditions for Attributes	
Available Written Capacity (Previously known as Effective Available Capacity)	The total amount of the provisioned capacity in the pools that is not allocated to the volumes in the pools.	
Capacity (Previously known as Pool Capacity)	The amount of space in the pools on the storage system that is available for creating volumes.	
Capacity-to-Limit	The amount of capacity that is available for storing data before the capacity limit is reached. For example, if you set a capacity limit, you can define a warning alert when the available capacity, in relation to the capacity limit, falls below the value that you specify such as:	
	Capacity-to-Limit ≤ 500 GiB	
Compression Savings	The estimated percentage of capacity that is saved by using data compression, across all pools on the storage system. The percentage is calculated across all compressed volumes in the pools and does not include the capacity of non-compressed volumes. Inline compression is a software feature that is supported by FlashSystem A9000 and FlashSystem A9000R, IBM Spectrum Accelerate, XIV storage systems with firmware version 11.6 or later, and resources that run IBM Spectrum Virtualize.	
Deduplication Savings	The estimated percentage of capacity that is saved by using data deduplication, across all data reduction pools on the storage system. The percentage is calculated across all deduplicated volumes in the pools and does not include the capacity of volumes that are not deduplicated. Available for FlashSystem A9000, FlashSystem A9000R, and resources that run IBM Spectrum Virtualize version 8.1.3 or later.	
File System Capacity	The total capacity on all of the file systems on the storage system or filer.	
Mapped Capacity (Previously known as Assigned Volume Space)	The total volume space in the storage system that is mapped or assigned to host systems, including child pool capacity.	
Overprovisioned Capacity (Previously known as Unallocatable Volume Space)	The capacity that cannot be allocated to volumes because the physical capacity of the pools cannot meet the demands for provisioned capacity. IBM Spectrum Control uses the following formula to determine this value: Provisioned Capacity — Capacity Available only for thin-provisioned volumes.	
Provisioned Capacity (Previously known as Total Volume Capacity)	The total storage space on all the volumes in pools. For thin-provisioned and compressed volumes, this value includes provisioned capacity. For volumes with parent pools, this value includes child pool capacity.	
Raw Capacity (Previously known as Raw Disk Capacity)	The total unformatted disk capacity of a storage system. When this value is calculated, IBM Spectrum Control does not include the capacity of storage system disks that become missing after data collection.	

Capacity Attributes	and conditions for capacity changes on storage systems (continued) Triggering Conditions for Attributes
Total Capacity Savings (Previously known as Total	The estimated percentage of capacity that is saved by using data deduplication, data compression, and thin provisioning.
Data Reduction Savings)	Available for FlashSystem A9000 and FlashSystem A9000R, IBM Spectrum Accelerate, XIV storage systems with firmware version 11.6 or later, and resources that run IBM Spectrum Virtualize.
Reserved Capacity (Previously known as Reserved Pool	The amount of unused capacity in the pool that is reserved for provisioning and optimization tasks.
Space)	Pool capacity is reserved when a provisioning or optimization task is created, and used when the task is run.
Shortfall	The difference between the amount of provisioned capacity that is committed to the volumes in the pools and the actual physical space that is available in the pools. As the provisioned capacity is allocated to the thin-provisioned and compressed volumes, the shortfall increases and becomes more critical.
	This value is determined by the formula, Overprovisioned Capacity ÷ Committed but Unused Capacity
	For example, the physical capacity of the pools is 70 GiB, but 150 GiB of provisioned capacity was committed to the thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB – 50 GiB) with only 20 GiB of available pool capacity (70 GiB – 50 GiB). Because only 20 GiB of the pool capacity is available, 80 GiB of the committed capacity cannot be allocated (100 GiB - 20 GiB).
Snapshot Space	The amount of space that is used by all of the snapshots of the file systems that are associated with the IBM Spectrum Scale cluster.
Unreserved Capacity (Previously known as Unreserved Pool Space)	The amount of space in storage system pools that is not allocated for volumes, and is not reserved by pending or scheduled provisioning tasks.
Unmapped Capacity	The total volume space in the storage system that is not mapped or assigned to host systems.
Unused Volume Capacity (Previously known as Effective Unallocated Volume Space)	The amount of the provisioned capacity in the storage pool that is not used.
Used Capacity (Previously known as Physical Allocation)	The percentage of physical capacity in pools that is allocated to volumes, including child pools. The value is always less than or equal to 100% because you cannot allocate more physical capacity to the volumes than is available in the pools. This value is determined by the formula, <i>Used Capacity ÷ Capacity × 100</i> . For example, if the capacity that is reserved for volumes is 50 GiB for a volume size of 200 GiB, used capacity is 25%.
Used Capacity (Previously known as Used Pool Space)	The capacity in the pool that is allocated to and used by volumes.
Used Written Capacity (%) (Previously known as Effective Used Capacity)	The percentage of capacity that is provisioned to the standard-provisioned volumes and the thin-provisioned volume, given the drive compression savings.

Table 43. Triggering attributes and conditions for capacity changes on storage systems (continued)		
Capacity Attributes	Triggering Conditions for Attributes	
Used Written Capacity (Previously known as Effective Used Capacity)	The total amount of provisioned capacity that is used by all volumes, given the drive compression savings.	
Written Capacity Limit (Previously known as Effective Capacity)	The amount of provisioned capacity that can be created, given the drive compression savings.	

Triggering conditions for storage system internal resource alerts

You can set up IBM Spectrum Control so that it examines the attributes, capacity, and performance of the internal resources of storage systems and notifies you when changes or violations are detected.

Alerts can notify you of general changes and performance issues on the following internal resources of storage systems:

- "Internal resources (common conditions)" on page 195
- "Clusters" on page 195
- "Disks" on page 195
- "Disk Groups" on page 196
- "FC Ports" on page 196
- "Filesets" on page 197
- "File Systems" on page 197
- "File System Pools" on page 197
- "Host Connections" on page 198
- "I/O Groups" on page 198
- "IP Ports" on page 199
- "Managed Disks" on page 199
- "Modules" on page 200
- "Network Shared Disks" on page 200
- "Nodes" on page 200
- "Pools" on page 203
- "Quotas" on page 207
- "RAID Arrays" on page 208
- "Shares" on page 208
- "Volumes" on page 209

Important: Not all the attributes upon which you can alert are listed here. To view a complete list of attributes upon which you can alert, go to **Settings** > **Alert Policies**. Double-click a default policy for a storage system. Click **Edit Alert Definitions** on the **Alert Definitions** tab. View the attributes that are available in the general, capacity, and performance categories. Note that the attributes that are automatically configured for alerts in the default alert policies, or default alerts, have a status of Active.

In the tables, default alerts are marked with an asterisk (*).

Tips:

• The type of storage system determines which attributes and performance conditions are available for alerts. For example, triggering conditions for shares are available only for storage systems that are configured for file storage, such as Storwize V7000 Unified.

• For capacity attributes, you can generate alerts when the amount of storage is greater than, less than, or equal to a specified value. You can also determine the unit of measurement for the attribute, such as KiB, MiB, GiB, or TiB.

Internal resources (common conditions)

There are a number of alert conditions that are common to many of the internal resources in a storage system. These common conditions represent key changes in your storage infrastructure. For example, you can specify conditions that generate alerts when specific internal resources are added to or deleted from a storage system, or when current data isn't being collected about resources.

By default, asset, capacity, and configuration metadata for storage systems is aggregated and collected daily. Define alerts to track daily changes to the attributes and conditions of the internal resources in your storage systems.

Table 44. Triggering attributes and conditions that are common to internal resources	
General Attributes	Triggering Conditions for Attributes
New resource	A resource is detected for the first time. Use this alert to be notified when new physical and logical resources are added to a storage system.
Removed <i>physical resource</i> Deleted <i>logical resource</i>	A previously monitored resource can no longer be found. Historical data about the resource is retained, but no current data is being collected. Use these alerts to be notified if a physical or logical resource is removed, deleted, or becomes unavailable.
	Physical resources include disks, RAID arrays, I/O groups, ports, nodes, host connections, clusters, and file systems. Logical resources include volumes, pools, filesets, and shares.

Clusters

Table 45. Triggering attributes and conditions for general changes on clusters		
General Attributes Triggering Conditions for Attributes		
New Cluster	A cluster is detected for the first time.	
Removed NAS Cluster	A previously monitored NAS cluster can no longer be found. Historical data about the cluster is retained, but no current data is being collected. Use this alert to be notified if a cluster is removed or becomes unavailable.	

Disks

Table 46. Triggering attributes and conditions for general changes on storage system disks	
General Attributes	Triggering Conditions for Attributes
New Disk	A disk is detected for the first time.
Removed Disk	A previously monitored disk can no longer be found. Historical data about the disk is retained, but no current data is being collected. Use this alert to be notified if a disk is removed or becomes unavailable.

Table 47. Triggering attributes and conditions for capacity changes on disks	
Capacity Attributes	Triggering Conditions for Attributes
Available Drive Capacity (Previously known as Available Disk Space)	The capacity that is available (not allocated) on the disk.
Capacity (Previously known as Disk Capacity)	The total amount of storage capacity that is on the disk.

Disk Groups

Table 48. Triggering attributes and conditions for general changes on storage system disk groups	
General Attributes	Triggering Conditions for Attributes
New Disk Group	A disk group is detected for the first time.
Removed Disk Group	A previously monitored disk group can no longer be found. Historical data about the disk group is retained, but no current data is being collected. Use this alert to be notified if a disk group is removed or becomes unavailable.
Status	One of the following statuses is detected for a disk group:
	Not Normal An error or warning condition is detected on a disk group.
	Warning A warning condition is detected on a disk group. This condition might occur if any of the disks in the disk group goes offline.
	Error An error condition is detected on a disk group. This condition might occur under the following conditions:
	 The percentage of remaining, unused volume capacity in a disk group that is not available to be used is too high.
	 The available capacity in a disk group that is not reserved for volumes is too low.
	The disk group provisioned capacity exceeds the warning or critical threshold boundary value.

FC Ports

Table 49. Triggering attributes and conditions for general changes on Fibre Channel ports	
General Attributes	Triggering Conditions for Attributes
New FC Port	A new FC port was detected for the first time.
Removed FC Port	A previously monitored FC port can no longer be found. Historical data about the port is retained, but no current data is being collected. Use this alert to be notified if a port is deleted or becomes unavailable.

Table 49. Triggering attributes and conditions for general changes on Fibre Channel ports (continued)	
General Attributes	Triggering Conditions for Attributes
Status	One of the following statuses is detected for a FC port:
	Not Normal An error or warning condition is detected on a port.
	Warning A warning condition is detected on a port.
	Error An error condition is detected on a port.

Filesets

Table 50. Triggering attributes and conditions for general changes on filesets	
General Attributes	Triggering Conditions for Attributes
Deleted NAS Fileset	A previously monitored NAS fileset can no longer be found. Historical data about the fileset is retained, but no current data is being collected. Use this alert to be notified if a fileset is deleted or becomes unavailable.
New Fileset	A fileset is detected for the first time.
State	A fileset is linked to a file system, unlinked from a file system, or deleting.

File Systems

Table 51. Triggering attributes and conditions for general changes on file systems	
General Attributes	Triggering Conditions for Attributes
New File System	A file system is detected for the first time.
Removed NAS File System	A previously monitored NAS file system can no longer be found. Historical data about the file system is retained, but no current data is being collected. Use this alert to be notified if a file system is deleted or becomes unavailable.

File System Pools

Table 52. Triggering attributes and conditions for general changes on file system pools	
General Attributes	Triggering Conditions for Attributes
Deleted Pool	A previously monitored NAS pool can no longer be found. Historical data about the pool is retained, but no current data is being collected. Use this alert to be notified if a pool is deleted or becomes unavailable.
New Pool	A NAS pool is detected for the first time.

Host Connections

Table 53. Triggering attributes and conditions for general changes on host connections	
General Attributes	Triggering Conditions for Attributes
New Host Connection	A host connection is detected for the first time.
Removed Host Connection	A previously monitored host connection can no longer be found. Historical data about the host connection is retained, but no current data is being collected. Use this alert to be notified if a host connection is removed or becomes unavailable.

Table 54. Triggering attributes and conditions for changes to performance metrics for unmap operations on host connections

Performance Attribute	Triggering Conditions for Attributes
Data Rate (Unmap)	Define an alert to monitor the average number of MiBs per second that were unmapped.
	This metric applies to systems that are running IBM Spectrum Virtualize V8.1.1 or later.
Overall I/O Rate (Unmap)	Define an alert to monitor the average number of unmap operations per second.
	This metric applies to systems that are running IBM Spectrum Virtualize V8.1.1 or later.
Peak Response Time (Unmap)	Define an alert to monitor the worst response time measured for an unmap operation in the sample interval.
	This metric applies to systems that are running IBM Spectrum Virtualize V8.1.1 or later.
Response Time (Unmap)	Define an alert to monitor the average number of milliseconds required to complete an unmap operation.
	This metric applies to systems that are running IBM Spectrum Virtualize V8.1.1 or later.
Unaligned Unmap I/O Rate	Define an alert to monitor the average number of volume unmap operations per second that are not aligned on an 8K boundary.
	This metric applies to systems that are running IBM Spectrum Virtualize V8.1.1 or later.

I/O Groups

Table 55. Triggering attributes and conditions for general changes on I/O groups	
General Attributes	Triggering Conditions for Attributes
New I/O Group	A new I/O group was detected for the first time.
Removed I/O Group	A previously monitored I/O group can no longer be found. Historical data about the I/O group is retained, but no current data is being collected. Use this alert to be notified if an I/O group is removed or becomes unavailable.

IP Ports

Triggering attributes and conditions are available for Internet Protocol ports on the nodes on block storage systems that run IBM Spectrum Virtualize, such as SAN Volume Controller, the IBM Storwize family, and some models of the IBM FlashSystem family.

Table 56. Triggering attributes and conditions for general changes on IP ports	
General Attributes	Triggering Conditions for Attributes
Status	One of the following statuses was detected for an IP port: Configured Unconfigured
Removed Port	A previously monitored IP port can no longer be found. Historical data about the port is retained, but no current data is being collected. Use this alert to be notified if a port is deleted or becomes unavailable.
New Port	A new IP port was detected for the first time.
Host Attach	An IP port was attached to, or detached from, a host.
Storage Attach	An IP port was attached to, or detached from, a storage system.
Management	One of the following management statuses was detected for an IP port: Configured Unconfigured
Remote Copy Relationship	The remote copy relationship changed, or one of the following remote copy statuses was detected for an IP port: • Active • Unconfigured • Standby

Managed Disks

Table 57. Triggering attributes and conditions for general changes on managed disks	
General Attributes	Triggering Conditions for Attributes
Managed Disk Status	One of the following statuses is detected for a managed disk: Not Normal An error or warning status was detected on a managed disk. Warning A warning status was detected on a managed disk. Error An error status was detected on a managed disk.
New Managed Disk	A new managed disk was detected for the first time.
Removed Managed Disk	A previously monitored managed disk can no longer be found. Historical data about the managed disk is retained, but no current data is being collected. Use this alert to be notified if a managed disk is removed or becomes unavailable.

Table 58. Triggering attributes and conditions for capacity changes on managed disks	
Capacity Attributes	Triggering Conditions for Attributes
Available Capacity	The unused storage capacity on the managed disk.
Capacity (Previously known as Total Space)	The total capacity on the managed disk on the storage system. This attribute is only available for Storwize V7000 storage systems that are configured as back-end storage.

Modules

Table 59. Triggering attributes and conditions for general changes on modules	
General Attributes	Triggering Conditions for Attributes
Deleted Module	A previously monitored module can no longer be found. Use this alert to be notified if a module is removed or becomes unavailable.

Network Shared Disks

Table 60. Triggering attributes and conditions for general changes on NSDs	
General Attributes	Triggering Conditions for Attributes
New NSD	A new NSD was detected for the first time.
Removed NSD	A previously monitored NSD can no longer be found. Historical data about the NSD is retained, but no current data is being collected. Use this alert to be notified if an NSD is removed or becomes unavailable.
Status	One of the following statuses is detected for an NSD: Not Normal An error or warning status was detected on an NSD. Warning A warning status was detected on an NSD. This status might occur if a storage system comes online or if its version changes. Error An error status was detected on an NSD.

Nodes

Table 61. Triggering attributes and conditions for general changes on nodes	
General Attributes	Triggering Conditions for Attributes
New Node	A new node was detected for the first time.
Removed Node (Block storage)	A previously monitored node can no longer be found. Historical data about the node is retained, but no current data is being collected. Use this alert to be notified if a node is removed or becomes unavailable.
Removed NAS Node (File storage)	A previously monitored NAS node can no longer be found. Historical data about the node is retained, but no current data is being collected. Use this alert to be notified if a node is removed or becomes unavailable.

Table 61. Triggering attributes and conditions for general changes on nodes (continued)	
General Attributes	Triggering Conditions for Attributes
Cloud Gateway Status (File storage)	One of the following statuses is detected for a node. This attribute is available only for IBM Spectrum Scale.
	Not Running The gateway service is not running because, for example, the service is stopped, no cloud account is configured, or the connection to the cloud provider failed.
	Not Installed The node is not a cloud gateway.
	No File System The node is a cloud gateway and the gateway service is running, but the node is not yet assigned to a file system.
	Stopped The node is a cloud gateway for a file system, but the gateway service is stopped.
	No Cloud Account The node is a cloud gateway for a file system and the gateway service is running, but no cloud account is configured.
	Disconnected The node is a cloud gateway, the gateway service is running, and a cloud account is configured, but the connection to the cloud provider failed.

Table 62. Triggering attributes of	Table 62. Triggering attributes and conditions for changes to performance metrics for cache fullness on nodes	
Performance Attributes	Triggering Conditions for Attributes	
Max Read Cache Fullness	Define an alert to monitor the maximum amount of the lower cache which the cache partitions of the pools that are managed by the node use for read operations. If the maximum value for the cache reaches 100%, the read cache partition for one or more of the pools is full. The read operations that pass through the node to the affected pools will be queued and the I/O response times will increase for the volumes in the affected pools.	
	This metric applies to systems that are running IBM Spectrum Virtualize V7.3 or later.	
Max Write Cache Fullness	Define an alert to monitor the maximum amount of the lower cache which the cache partitions of the pools that are managed by the node use for write operations. If the maximum value for the cache reaches 100%, the write cache partition for one or more of the pools is full. The write operations that pass through the node to the affected pools will be queued and the I/O response times will increase for the volumes in the affected pools.	
	This metric applies to systems that are running IBM Spectrum Virtualize V7.3 or later.	
Read Cache Fullness	Define an alert to monitor the average amount of the lower cache which the cache partitions of the pools that are managed by the node use for read operations. Use this alert to monitor the average cache fullness for read operations to identify the nodes that experience heavy cache usage.	
	This metric applies to systems that are running IBM Spectrum Virtualize V7.3 or later.	

Table 62. Triggering attributes and conditions for changes to performance metrics for cache fullness on nodes (continued)

Performance Attributes	Triggering Conditions for Attributes
Write Cache Fullness	Define an alert to monitor the average amount of the lower cache which the cache partitions of the pools that are managed by the node are using for write operations. Use this alert to monitor the average cache fullness for write operations to identify the nodes that experience heavy cache usage.
	This metric applies to systems that are running IBM Spectrum Virtualize V7.3 or later.

Table 63. Triggering attributes and conditions for changes to performance metrics for recovering data in data reduction pools on nodes

Performance Attributes	Triggering Conditions for Attributes
Data Rewrite Rate	Define an alert to monitor the rate at which data is rewritten when a host overwrites data in data reduction pools on the node.
	This metric applies to systems that are running IBM Spectrum Virtualize V8.1.2 or later.
Extent Collection Rate	Define an alert to monitor the number of volume extents that were processed for garbage collection. The reclaimable capacity in the volume extents is collected so that it can be reused in the data reduction pools on the node.
	This metric applies to systems that are running IBM Spectrum Virtualize V8.1.2 or later.
Data Movement Rate	Define an alert to monitor the rate at which valid data in a reclaimed volume extent is moved to a new extent in the data reduction pool on the node.
	This metric applies to systems that are running IBM Spectrum Virtualize V8.1.2 or later.
New Address Write Rate	Define an alert to monitor the rate at which capacity is used to write the host's data to unallocated addresses in the data reduction pool on the node. Use this alert to determine which hosts are increasing the amount of capacity that is being written to data reduction pools on a node.
	This metric applies to systems that are running IBM Spectrum Virtualize V8.1.2 or later.
Reclaimable Capacity	Define an alert to monitor the amount of capacity that can be reclaimed in the data reduction pools on the node.
	This metric applies to systems that are running IBM Spectrum Virtualize V8.1.2 or later.
Recovered Capacity Rate	Define an alert to monitor the rate at which capacity is recovered by garbage collection for reuse in the data reduction pools on the node.
	This metric applies to systems that are running IBM Spectrum Virtualize V8.1.2 or later.

Pools

Table 64. Triggering attributes and conditions for general changes on pools	
General Attributes	Triggering Conditions for Attributes
Deleted Pool	A previously monitored pool can no longer be found. Historical data about the pool is retained, but no current data is being collected. Use this alert to be notified if a pool is deleted or becomes unavailable.
New Storage Pool	A new pool was detected for the first time.
Status	One of the following statuses is detected for a pool:
	Not Normal An error or warning condition is detected on a pool. Warning
	A warning condition is detected on a pool. This condition might occur if a pool goes offline.
	Error An error condition is detected on a pool. This condition might occur under the following conditions:
	The percentage of remaining, unused volume capacity in a pool that is not available to be used is too high.
	The available capacity in a pool that is not reserved for volumes is too low.
	The pool provisioned capacity exceeds the warning or critical threshold boundary value.

Table 65. Triggering attributes and conditions for capacity changes on pools	
Capacity Attributes	Triggering Conditions for Attributes
Adjusted Used Capacity	The amount of capacity that can be used without exceeding the capacity limit. For example, you set a capacity limit of 80% for your pools. You want to get an informational alert when the adjusted used capacity exceeds 60% and a critical alert when the adjusted used capacity exceeds 80%. So, you define an informational alert with these parameters:
	Adjusted Used Capacity ≥ 60%
	And, you define a critical alert with these parameters:
	Adjusted Used Capacity ≥ 80%
Available Repository Capacity	The amount of available, unallocated storage space on all extents in the repository of a pool for Track Space-Efficient (TSE) thin-provisioning. This attribute applies only to the DS8000 storage systems. You can use this alert to be notified about space-efficient volumes.
Available Virtual Capacity	The amount of provisioned capacity in a thin-provisioned pool that is not used by volumes.
Available Written Capacity (Previously known as Effective Available Capacity)	The total amount of the provisioned capacity in the pools that is not allocated to the volumes in the pools.

Capacity Attributes	Triggering Conditions for Attributes
Capacity-to-Limit	The amount of capacity that is available for storing data before the capacity limit is reached. For example, if you set a capacity limit, you can define a warning alert when the available capacity relative to the capacity limit, falls below the value that you specify, such as:
	Capacity-to-Limit ≤ 500 GiB
Compression Savings	The estimated percentage of capacity that is saved by using data compression. The percentage is calculated across all compressed volumes in the pool and does not include the capacity of non-compressed volumes. Inline compression is a software feature that is supported by FlashSystem A9000 and FlashSystem A9000R, IBM Spectrum Accelerate, XIV storage systems with firmware version 11.6 or later, and resources that run IBM Spectrum Virtualize.
Deduplication Savings	The estimated percentage of capacity that is saved by using data deduplication. The percentage is calculated across all deduplicated volumes in the pool and does not include the capacity of volumes that are not deduplicated. Available for resources that run IBM Spectrum Virtualize version 8.1.3 or later.
Mapped Capacity (Previously known as Assigned Volume	The amount of space on all the volumes in a pool that are mapped or assigned to host systems.
Space)	For a thin-provisioning pool, this value includes the provisioned capacity of thin-provisioned volumes, which might exceed the total capacity in the pool. For Hitachi VSP non-thin provisioning pool capacity, this value is the sum of assigned regular host-accessible volumes. Volumes that are used for thin-provisioning (pool volumes) are not included.
Enterprise HDD Capacity	The total capacity on Enterprise hard disk drives in the pool. Easy Tier can use the capacity on these drives to retier the volume extents in the pool.
Enterprise HDD Available Capacity	The available capacity on Enterprise hard disk drives in the pool. Easy Tier can use the available capacity to retier the volume extents in the pool.
Enterprise HDD Available Capacity (%)	The percentage of capacity on Enterprise hard disk drives in the pool that is available. Easy Tier can use the available capacity to retier the volume extents in the pool.
Nearline HDD Capacity	The total capacity on Nearline hard disk drives in the pool. Easy Tier can use the capacity on these drives to retier the volume extents in the pool.
Nearline HDD Available Capacity	The available capacity on Nearline hard disk drives in the pool. Easy Tier can use the available capacity to retier the volume extents in the pool.
Nearline HDD Available Capacity (%)	The percentage of capacity on Nearline hard disk drives in the pool that is available. Easy Tier can use the available capacity to retier the volume extents in the pool.
Provisioned Capacity (Previously known as Virtual Allocation)	The percentage of physical space in a pool that is committed to the provisioned capacity of the volumes in the pool.
	You can use this alert to be notified about space-efficient volumes.
Repository Space	The amount of space on all extents in the repository of a pool. This attribute applies only to the DS8000 storage systems.

	and conditions for capacity changes on pools (continued)
Capacity Attributes	Triggering Conditions for Attributes
Reserved Capacity (Previously known as Reserved Pool Space)	The amount of unused capacity in a pool that is reserved for provisioning and optimization tasks.
	Pool capacity is reserved when a provisioning or optimization task is created, and used when the task is run.
Reserved Volume Capacity (Previously known as Unused Space)	The amount of pool capacity that is reserved but has not been used yet to store data on the thin-provisioned volume.
Shortfall	The difference between the amount of provisioned capacity that is committed to the volumes in the pools and the actual physical capacity that is available in the pools. As the provisioned capacity is allocated to the thin-provisioned and compressed volumes, the shortfall increases and becomes more critical.
	This value is determined by the formula, Overprovisioned Capacity ÷ Committed but Unused Capacity
	For example, the physical capacity of the pools is 70 GiB, but 150 GiB of provisioned capacity was committed to the thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB – 50 GiB) with only 20 GiB of available pool capacity (70 GiB – 50 GiB). Because only 20 GiB of the pool capacity is available, 80 GiB of the committed capacity cannot be allocated (100 GiB - 20 GiB).
SCM Capacity	The total capacity on Storage Class Memory (SCM) drives in the pool. Easy Tier can use the capacity on SCM drives to retier the volume extents in the pool.
	Available for IBM Spectrum Virtualize systems, such as IBM FlashSystem 9100, IBM FlashSystem 7200, and the IBM Storwize family.
SCM Available Capacity	The available capacity on Storage Class Memory (SCM) drives in the pool. Easy Tier can use the available capacity to retier the volume extents in the pool.
	Available for IBM Spectrum Virtualize systems, such as IBM FlashSystem 9100, IBM FlashSystem 7200, and the IBM Storwize family.
SCM Available Capacity (%)	The percentage of capacity on Storage Class Memory (SCM) drives in the pool that is available. Easy Tier can use the available capacity to retier the volume extents in the pool.
	Available for IBM Spectrum Virtualize systems, such as IBM FlashSystem 9100, IBM FlashSystem 7200, and the IBM Storwize family.
Tier 0 Flash Capacity	The total capacity on Tier 0 flash solid-state drives in the pool. Easy Tier can use the capacity on these drives to retier the volume extents in the pool.
Tier 0 Flash Available Capacity	The available capacity on Tier 0 flash solid-state drives in the pool. Easy Tier can use the available capacity to retier the volume extents in the pool.
Tier 0 Flash Available Capacity (%)	The percentage of capacity on Tier 0 flash solid-state drives in the pool that is available. Easy Tier can use the available capacity to retier the volume extents in the pool.
Tier 1 Flash Capacity	The total capacity on Tier 1 flash solid-state drives in the pool. Easy Tier can use the capacity on these drives to retier the volume extents in the pool.
Tier 1 Flash Available Capacity	The available capacity on Tier 1 flash solid-state drives in the pool. Easy Tier can use the available capacity to retier the volume extents in the pool.

Capacity Attributes	Triggering Conditions for Attributes
Tier 1 Flash Available Capacity (%)	The percentage of capacity on Tier 1 flash solid-state drives in the pool that is available. Easy Tier can use the available capacity to retier the volume extents in the pool.
Tier 2 Flash Capacity	The total capacity on Tier 2 flash, high-capacity drives in the pool. Easy Tier can use the capacity on these drives to retier the volume extents in the pool. Available for DS8000 storage systems.
Tier 2 Flash Available Capacity	The available capacity on Tier 2 flash, high-capacity drives in the pool. Easy Tier can use the available capacity to retier the volume extents in the pool. Available for DS8000 storage systems.
Tier 2 Flash Available Capacity (%)	The percentage of capacity on Tier 2 flash, high-capacity drives in the pool that is available. Easy Tier can use the available capacity to retier the volume extents in the pool.
	Available for DS8000 storage systems.
Total Capacity Savings (Previously known as Total	The estimated percentage of capacity that is saved by using data deduplication, data compression, and thin provisioning.
Data Reduction Savings)	Available for FlashSystem A9000 and FlashSystem A9000R, IBM Spectrum Accelerate, XIV storage systems with firmware version 11.6 or later, and resources that run IBM Spectrum Virtualize.
Total Reserved Capacity	The total amount of space on the pool that is reserved for provisioning and optimization tasks. Pool space is reserved when a provisioning or optimization task is created, and allocated when the task is run.
Used Capacity (Previously known as Allocated Space)	The amount of space that is reserved for all the volumes in a pool.
Used Capacity (%) (Previously known as Physical Allocation)	The percentage of physical capacity in the pool that is used by the standard-provisioned volumes, the thin-provisioned volumes, and the volumes in child pools. This value is always less than or equal to 100% because you cannot allocate more physical space than is available in a pool.
	This value is determined by the formula, <i>Used Capacity ÷ Capacity × 100</i> . For example, if the space that is reserved for volumes is 50 GiB for a volume size of 200 GiB, used capacity is 25%.
Used Repository Space	The amount of used capacity on all extents in the repository of a pool. This attribute applies only to the DS8000 storage systems.
	You can use this alert to be notified about space-efficient volumes.
Used Volume Space	The amount of space on the storage system that is used by volumes.
Used Written Capacity (%) (Previously known as Effective Used Capacity)	The percentage of capacity that is provisioned to the standard-provisioned volumes and the thin-provisioned volumes, given the drive compression savings.
Used Written Capacity (Previously known as Effective Used Capacity)	The total amount of provisioned capacity that is used by all the volumes given the drive compression savings.

Table 65. Triggering attributes and conditions for capacity changes on pools (continued)	
Capacity Attributes	Triggering Conditions for Attributes
User Reserved Capacity	The amount of space in the pools on the storage system that is reserved for user-defined purposes.
Virtual Capacity Limit	The maximum amount of virtual storage space available to allocate to volumes in the storage pools that are associated with the storage system. You can use this alert to be notified about space-efficient volumes.
Virtual Volume Space	The total amount of physical space in a pool that is committed to the total virtual capacity of the volumes in the pool. You can use this alert to be notified about space-efficient volumes.
Written Capacity Limit (Previously known as Effective Capacity)	The amount of provisioned capacity that can be created, given the drive compression savings.

Table 66. Triggering attributes and conditions for changes to performance metrics for cache fullness on pools	
Performance Attributes	Triggering Conditions for Attributes
Read and Write Cache Fullness	Define an alert to monitor the average amount of the lower cache which the pools' cache partitions on the nodes use for read and write operations. Use this alert to monitor the average cache fullness for read and write operations to identify the pools that experience heavy cache usage.
	This metric applies to systems that are running IBM Spectrum Virtualize V7.3 or later.
Max Read and Write Cache Fullness	Define an alert to monitor the maximum amount of the lower cache which the cache partitions on the nodes that manage the pool use for read and write operations. If the maximum value for the cache reaches 100%, one or more cache partitions on one or more pools is full. The operations that pass through the pools with full cache partitions will be queued and I/O response times will increase for the volumes in the affected pools. This metric applies to systems that are running IBM Spectrum Virtualize V7.3 or later.

Quotas

Table 67. Triggering attributes and conditions for general changes on quotas	
General Attributes	Triggering Conditions for Attributes
Deleted Quota	A previously monitored quota can no longer be found. Historical data about the quota is retained, but no current data is being collected. Use this alert to be notified if a quota is deleted or becomes unavailable.
New Quota	A new quota was detected for the first time.

RAID Arrays

Table 68. Triggering attributes and conditions for general changes on RAID arrays	
General Attributes	Triggering Conditions for Attributes
New RAID Array	A new RAID array was detected for the first time.
Removed RAID Array	A previously monitored RAID array can no longer be found. Historical data about the RAID array is retained, but no current data is being collected. Use this alert to be notified if a RAID array is removed or becomes unavailable.
Status	One of the following statuses is detected for a RAID array: Not Normal An error or warning condition is detected on a RAID array. Warning A warning condition is detected on a RAID array. Error An error condition is detected on a RAID array.

Table 69. Triggering attributes and conditions for capacity changes on RAID arrays	
Capacity Attributes	Triggering Conditions for Attributes
Available Physical Space	The amount and percentage of storage space that is unused in the RAID array. Available for RAID arrays with disk drive modules that use inline data compression, such as RAID arrays on FlashSystem 9100 and FlashSystem 900.
RAIDArrayFreeSpace	The amount of unused storage space in the RAID array.
Capacity (Previously known as Total Space)	The total capacity in the RAID array.

Shares

Table 70. Triggering attributes and conditions for general changes on shares	
General Attributes	Triggering Conditions for Attributes
New Share	A new share was detected for the first time.
Deleted Share	A previously monitored share can no longer be found. Historical data about the share is retained, but no current data is being collected. Use this alert to be notified if a share is deleted or becomes unavailable.

Volumes

Table 71. Triggering attributes and conditions for general changes on volumes	
General Attributes	Triggering Conditions for Attributes
Deleted Volume	A previously monitored volume can no longer be found. Historical data about the volume is retained, but no current data is being collected. Use this alert to be notified if a volume is deleted or becomes unavailable.
	If you are doing tasks where many volumes are being deleted, you might want to temporarily disable alerts that use the Deleted Volume attribute. For example, you might want to disable Deleted Volume alerts temporarily if you are doing maintenance tasks or decommissioning storage.
New Volume	A new volume was detected for the first time.
Status	One of the following statuses is detected for a volume:
	Not Normal An error or warning condition is detected on a RAID array. Warning A warning condition is detected on a RAID array. Error
	An error condition is detected on a RAID array.

Table 72. Triggering attributes and conditions for capacity changes on volumes	
Capacity Attributes	Triggering Conditions for Attributes
Compression Savings	The estimated percentage of capacity that is saved by using data compression. Inline compression is a software feature that is supported by FlashSystem A9000 and FlashSystem A9000R, IBM Spectrum Accelerate, XIV storage systems with firmware version 11.6 or later, and resources that run IBM Spectrum Virtualize.
Provisioned Capacity (Previously known as Unallocatable Space)	The amount of space by which the capacity of a volume exceeds the physical capacity of the associated pool. In thin-provisioned environments, it is possible to over commit (over provision) storage in a pool by creating volumes with more virtual capacity than can be physically allocated in the pool. This value represents the amount of volume space that cannot be allocated based on the current capacity of the pool.
Real Capacity	The total amount of storage space that is physically allocated to a volume. For thin-provisioned volumes, this value is less than the provisioned capacity of the volume. In an XIV and IBM Spectrum Accelerate, this value represents the physical ("hard") capacity of the volume, not the provisioned ("soft") capacity.
Remaining Managed Space	The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.
Reserved Volume Capacity (Previously known as Unused Space)	The amount of pool capacity that is reserved but has not been used yet to store data on the thin-provisioned volume.
	The value for Reserved Volume Capacity is available only for SAN Volume Controller and Storwize family storage systems that are configured with block storage.

Table 72. Triggering attributes and conditions for capacity changes on volumes (continued)	
Capacity Attributes	Triggering Conditions for Attributes
Space	The amount of space in a pool that is allocated to a volume.
Unused Capacity (Previously known as Unallocated Space)	The capacity in a pool that is not reserved for a volume. This value is determined by the formula: Capacity – Used Capacity The value for Unused Capacity is available only for thin provisioned volumes.
Uncompressed Used Capacity	The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.
Used Capacity (Previously known as Allocated Space)	The capacity on a pool that is physically allocated to a volume.
Written Capacity (Previously known as Written Space)	The amount of data that is written from the assigned hosts to the volume before compression or data deduplication are used to reduced the size of the data.

Table 73. Triggering attributes and conditions for changes to performance metrics for unmap operations on volumes **Performance Attribute Triggering Conditions for Attributes** Data Rate (Unmap) Define an alert to monitor the average number of MiBs per second that were unmapped. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.1 or later. Overall I/O Rate (Unmap) Define an alert to monitor the average number of unmap operations per second. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.1 or later. Peak Response Time (Unmap) Define an alert to monitor the worst response time measured for an unmap operation in the sample interval. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.1 or later. Response Time (Unmap) Define an alert to monitor the average number of milliseconds required to complete an unmap operation. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.1 or later. Unaligned Unmap I/O Rate Define an alert to monitor the average number of volume unmap operations per second that are not aligned on an 8K boundary. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.1 or later.

Triggering conditions for hypervisor alerts

You can set up IBM Spectrum Control so that it examines the attributes and capacity of a hypervisor and notifies you when changes are detected.

Alerts can notify you of general changes and capacity changes on the following resources:

- "Hypervisors" on page 211
- "Controllers" on page 212
- "Data Stores" on page 212
- "Disks" on page 213
- "Paths" on page 214
- "Virtual Machines" on page 214
- "VMDKs" on page 215

Important: Not all the attributes upon which you can alert are listed here. A number of other attributes are available for alerts and are based on the key properties of a hypervisor. To view a complete list of hypervisor attributes upon which you can alert, go to **Settings** > **Alert Policies**. Double-click a default policy for a hypervisor. Click **Edit Alert Definitions** on the **Alert Definitions** tab. View the attributes that are available in the general and capacity categories. Note that the attributes that are automatically configured for alerts in the default alert policies, or default alerts, have a status of Active.

In the tables, default alerts are marked with an asterisk (*).

Tip: For capacity attributes, you can generate alerts when the amount of storage is greater than, less than, or equal to a specified value. You can also determine the unit of measurement for the attribute, such as KiB, MiB, GiB, or TiB.

Hypervisors

Table 74. Triggering attributes and conditions for general hypervisor changes	
General Attributes	Descriptions of Triggering Conditions
Last Successful Probe	A specified amount of time has passed since a probe was able to collect data about a hypervisor. You can use this alert to be notified when up-to-date configuration and status data is not being collected about a hypervisor and its existing data might be stale. Data collection might be interrupted or not occur if the resource, network, or IBM Spectrum Control server is unavailable.
Probe Status*	One of the following statuses is detected for a probe:
	Not Successful An error or warning occurred during data collection. This status indicates that a probe did not collect any data, or only collected a partial set of data about a resource.
	Warning A probe completed, but might not have collected a complete set of data. This status might occur if data cannot be collected about one or more of the internal resources of a resource.
	Error (default) A probe did not complete when it attempted to collect asset data about the resource. This status might occur if the resource cannot be reached during data collection.
	For details about why a specific status occurred, check the log for the probe. To check the log, go to the details page for a resource, click Data Collection , and select Actions > Open Logs in the Probe section on the Data Collection page.

Table 74. Triggering attributes and conditions for general hypervisor changes (continued)	
General Attributes	Descriptions of Triggering Conditions
Status	One of the following statuses is detected on a hypervisor:
	Not Normal An error or warning status was detected on the hypervisor or its internal resources.
	Warning A warning status was detected on the hypervisor or its internal resources.
	Error (default) An error status was detected on the hypervisor or its internal resources. For example, an error status might occur when a hypervisor goes offline, or a server disk is disconnected or partially disconnected if it has multiple paths and one of the paths is disconnected.
	Unreachable One or more of the monitored resources for a hypervisor are not responding. This status might be caused by a problem in the network.

Controllers

Table 75. Triggering attributes and conditions for general changes on controllers	
Controller Attributes	Descriptions of Triggering Conditions
New Disk Controller	A disk controller is detected for the first time. Use this alert to be notified of hardware additions on servers.
Removed Disk Controller	A previously monitored disk controller can no longer be found. Historical data about the controller is retained, but no current data is being collected. Use this alert to be notified if a controller is removed or becomes unavailable.

Data Stores

Table 76. Triggering attributes and conditions for general changes on data stores	
Data Store Attributes	Descriptions of Triggering Conditions
Deleted Logical Volume	A previously monitored logical volume on a data store can no longer be found. Historical data about the volume is retained, but no current data is being collected. Use this alert to be notified if a volume is deleted or becomes unavailable.
New Logical Volume	A logical volume on the data store is detected for the first time.
VMDKs	You can use a number of operators to determine when you are notified of the number of virtual machine disk files (VMDKs) on the data store. For example, you can select to be notified when the number of VMDKs changes, or when the number is or is not a specific number, or when the number lies outside a specific range.

Table 77. Triggering attributes and conditions for capacity changes on data stores	
Capacity Attributes	Descriptions of Triggering Conditions
Available Data Store Space	The amount of unused storage capacity on the hypervisor data store.

Table 77. Triggering attributes and conditions for capacity changes on data stores (continued)	
Capacity Attributes	Descriptions of Triggering Conditions
Available File System Capacity	The total amount of unused storage capacity on a file system.
Data Store Capacity	The total amount of storage capacity that is assigned to a data store.
File System Capacity (Previously known as Total File System Capacity)	The capacity of a file system.
Used File System Capacity	The used capacity on the hypervisor file system.
Used Capacity	The percentage of storage capacity that is used on the hypervisor data store.

Disks

Table 78. Triggering attributes and conditions for general changes on disks	
General Attributes	Descriptions of Triggering Conditions
Firmware	The version of the Licensed Internal Code on the disk changes. You can use a number of operators to determine when you are notified of a firmware change, such as when the firmware is, or is not, a specific version, or when the version number contains a specific value.
Multipathing Policy	The multipathing policy that is in effect for a disk. For example, you can be notified when the policy changes, or when the policy is Round Robin, Load Balancing, Failover Only, or other another policy.
New Disk	A disk is detected for the first time. Use this alert to be notified of hardware changes on servers or hypervisors.
Paths	The number of access paths that are associated with the disk falls outside a specified range, or is equal to or not equal to a specified value.
Removed Disk	A previously monitored disk can no longer be found. Historical data about the disk is retained, but no current data is being collected. Use this alert to be notified if a disk is removed or becomes unavailable.
Status*	One of the following statuses is detected on a disk: Not Normal An error or warning status was detected on the disk. Warning A warning status was detected on the disk. Error (default) An error status was detected on the disk.

Table 79. Triggering attributes and conditions for capacity changes on disks	
Capacity Attributes	Descriptions of Triggering Conditions
Available Drive Capacity (Previously known as Available Disk Space)	The unused capacity on a hypervisor disk.
Capacity	The total amount of storage capacity assigned to a hypervisor disk.
Used Capacity	The amount of used storage capacity on a hypervisor disk .

Paths

Table 80. Triggering attributes and conditions for general changes on paths	
Path Attributes	Descriptions of Triggering Conditions
Deleted Path	A previously monitored access path for a server disk can no longer be found. This change might or might not affect the availability of the disk because there might be more than one path available.
New Path	An access path for a disk is detected for the first time.
Status*	One of the following statuses is detected on a path: Not Normal An error or warning status was detected on the path. Warning A warning status was detected on the path. Error (default) An error status was detected on the path.

Virtual Machines

Table 81. Triggering attributes and conditions for general changes on virtual machines	
Virtual Machine Attributes	Descriptions of Triggering Conditions
Deleted Virtual Machine	A previously monitored access path for a virtual machine can no longer be found. Historical data about the virtual machine is retained, but no current data is being collected. Use this alert to be notified if a virtual machine is deleted or becomes unavailable.
New Virtual Machine	A virtual machine is detected for the first time.
Status	One of the following statuses is detected on a virtual machine: Not Normal An error or warning status was detected on the virtual machine. Warning A warning status was detected on the virtual machine. Error (default) An error status was detected on the virtual machine.

Table 81. Triggering attributes and conditions for general changes on virtual machines (continued)	
Virtual Machine Attributes	Descriptions of Triggering Conditions
VMDKs	You can use a number of operators to determine when you are notified of the number of VMDKs on the virtual machine. For example, you can select to be notified when the number of VMDKs changes, or when the number is or is not a specific number, or when the number lies outside a specific range.

Table 82. Triggering attributes and conditions for capacity changes on virtual machines	
Capacity Attributes	Descriptions of Triggering Conditions
Capacity	The total amount of storage capacity assigned to a virtual machine hosted by the hypervisor.

VMDKs

Table 83. Triggering attributes and conditions for general changes on virtual machine disks (VMDKs)	
VMDK Attributes	Descriptions of Triggering Conditions
Deleted Virtual Disk	A previously monitored virtual disk can no longer be found. Historical data about the virtual disk is retained, but no current data is being collected. Use this alert to be notified if a virtual disk is deleted or becomes unavailable.
New Virtual Disk	A virtual disk is detected for the first time.

Table 84. Triggering attributes and conditions for capacity changes on VMDKs		
Capacity Attributes	Descriptions of Triggering Conditions	
Size	The size of the VMDK on the hypervisor data store.	
Used Capacity	The percentage of storage capacity that is used on the VMDK on the hypervisor data store. The used capacity is available if a Storage Resource agent is deployed.	

Triggering conditions for switch alerts

You can set up IBM Spectrum Control so that it examines the attributes and performance of a switch and notifies you when changes or violations are detected.

Alerts can notify you of general changes and performance issues on the following resources:

- Switches (performance)
- "Switches (general changes)" on page 216
- "Ports" on page 218

Important: Not all the attributes upon which you can alert are listed here. A number of other attributes are available for alerts and are based on the key properties of a switch. To view a complete list of attributes upon which you can alert, go to Settings > Alert Policies. Double-click a default policy for a switch. Click Edit Alert Definitions on the Alert Definitions tab. View the attributes that are available in the general, capacity, and performance categories. Note that the attributes that are automatically configured for alerts in the default alert policies, or default alerts, have a status of Active.

In the tables, default alerts are marked with an asterisk (*).

Switches (performance)

Define alerts that notify you when the performance of a switch falls outside a specified threshold. In alerts, you can specify conditions based on metrics that measure the performance of switch ports, including I/O, data, and error rates, and frame transfer sizes. By creating alerts with performance conditions, you can be informed about potential bottlenecks in your network infrastructure.

For example, you can define an alert to be notified when the port congestion index for a port is greater than or equal to a specified threshold. Port congestion represents the estimated degree to which frame transmission was delayed due to a lack of buffer credits. Use this alert to help identify port conditions that might slow the performance of the resources to which those ports are connected.

You can also be notified when a metric is less than a specified threshold, such as when you want to identify ports that might be under used.

For a complete list of switch metrics that can be alerted upon, see <u>"Performance metrics for switches" on page 343.</u>

Tips for performance conditions:

- A performance monitor must collect data about a resource before IBM Spectrum Control can determine whether a threshold is violated and an alert is generated for a performance condition.
- When you define a performance alert for the inter-switch connections category of a switch, the performance of ISL trunk connections is measured and alerted on. To alert on the performance of standard ISL connections, define performance alerts for the ports category of a switch.

Best practice: When you set thresholds for performance conditions, try to determine the best value so you can derive the maximum benefit without generating too many false alerts. Because suitable thresholds are highly dependent on the type of workload that is being run, hardware configuration, the number of physical disks, exact model numbers, and other factors, there are no easy or standard default rules.

A recommended approach is to monitor the performance of resources for a number of weeks and by using this historical data, determine reasonable threshold values for each performance condition. After that is done, you can fine-tune the condition settings to minimize the number of false alerts.

Switches (general changes)

Table 85. Triggering attributes and conditions for general changes on switches		
General Attributes	Triggering Conditions for Attributes	
Firmware	The version of the Licensed Internal Code on the switch changes. You can use a number of operators to determine when you are notified of a firmware change, such as when the firmware is, or is not, a specific version, or when the version number contains a specific value.	
Last Successful Probe Last Successful Monitor	A specified amount of time has passed since a probe or performance monitor was able to collect data about a switch. You can use this alert to be notified when up-to-date configuration, status, or performance data is not being collected about a switch and its existing data might be stale. This situation might occur if the resource, network, or IBM Spectrum Control server is unavailable.	

	s and conditions for general changes on switches (continued)
General Attributes	Triggering Conditions for Attributes
Probe Status*	One of the following statuses is detected for a probe:
	Not Successful An error or warning occurred during data collection. This status indicates that a probe did not collect any data, or only collected a partial set of data about a resource.
	Warning A probe completed, but might not have collected a complete set of data. This status might occur if data cannot be collected about one or more of the internal resources of a resource.
	Error (default) A probe did not complete when it attempted to collect asset data about the resource. This status might occur if the resource cannot be reached during data collection.
	For details about why a specific status occurred, check the log for the probe. To check the log, go to the details page for a resource, click Data Collection , and select Actions > Open Logs in the Probe section on the Data Collection page.
Performance Monitor Status*	One of the following statuses is detected for a performance monitor:
	Not Normal An error or warning occurred during data collection. This status indicates that a performance monitor did not collect any data, or only collected a partial set of data about a resource.
	Warning A performance monitor completed, but did not collect a complete set of performance data. This status might occur if the resource was rebooted during data collection, no valid performance data was provided by the resource, or a communication error occurred with the resource or its associated agent.
	Error A performance did not complete when it attempted to collect performance data about the resource. This status might occur if the resource cannot be reached during data collection, or if no configuration data is available for the resource.
	For details about why a specific status occurred, check the log for the performance monitor. To check the log, go to the details page for a resource, click Data Collection , and select Actions > Open Logs in the Performance Monitor section on the Data Collection page.

Table 85. Triggering attributes and conditions for general changes on switches (continued)	
General Attributes	Triggering Conditions for Attributes
Status*	One of the following conditions is detected on a switch:
	Not Normal An error or warning status was detected on the switch or its internal resources.
	Warning A warning status was detected on the switch or its internal resources.
	Error (default) An error status was detected on the switch or its internal resources. For example, an error status might occur when a switch goes offline.
	Unreachable One or more of the monitored resources for a switch are not responding. This status might be caused by a problem in the network.

Ports

Table 86. Triggering attributes and conditions for general changes on ports		
General Attributes	Triggering Conditions for Attributes	
Removed Port	A previously monitored port can no longer be found. Historical data about the port is retained, but no current data is being collected. Use this alert to be notified if a port is removed or becomes unavailable.	
Speed	The negotiated speed of a port, which is measured in GiB per second. You can generate an alert when the port speed changes, or when it is greater than, less than, equal to, or is not equal to a specific value.	
State	A port is online, enabled but offline, or disabled.	
Status*	One of the following statuses is detected for a port:	
	Not Normal An error or warning status is detected on a port.	
	Warning A warning status is detected on a port. This status might occur if the switch is stopped, starting, or in service (being maintained, cleaned, or administered).	
	Error An error status is detected on a port.	

Triggering conditions for fabric alerts

You can set up IBM Spectrum Control so that it examines the attributes of a fabric and notifies you when changes are detected.

Alerts can notify you of general changes on fabric resources.

Important: Not all the attributes upon which you can alert are listed here. A number of other attributes are available for alerts and are based on the key properties of fabrics or switches. To view a complete list of attributes upon which you can alert, go to Settings > Alert Policies. Double-click a default policy for a fabric or switch. Click Edit Alert Definitions on the Alert Definitions tab. View the attributes that are available in the general and capacity categories, and in the performance category for switches. Note that the attributes that are automatically configured for alerts in the default alert policies, or default alerts, have a status of Active.

Not all attributes are available for all types of fabrics. For example, you can't probe fabrics that contain Cisco switches or Brocade switches that run Fabric OS 8.2.1 or later. Therefore the Last Successful Probe and Probe Status attributes are not available for those types of fabric.

In the table, default alerts are marked with an asterisk (*).

Fabrics

Table 87. Triggering attributes and conditions for general changes on fabrics	
Fabric Attributes	Triggering Conditions for Attributes
Last Successful Probe	A specified amount of time has passed since a probe was able to collect data about a fabric. You can use this alert to be notified when up-to-date configuration and status data is not being collected about a fabric and its existing data might be stale. This situation might occur if the fabric, network, or IBM Spectrum Control server is unavailable.
Probe Status*	One of the following statuses is detected for a probe:
	Not Successful An error or warning occurred during data collection. This status indicates that a probe did not collect any data, or only collected a partial set of data about a resource.
	Warning A probe completed, but might not have collected a complete set of data. This status might occur if data cannot be collected about one or more of the internal resources of a resource.
	Error (default) A probe did not complete when it attempted to collect asset data about the resource. This status might occur if the resource cannot be reached during data collection.
	For details about why a specific status occurred, check the log for the probe. To check the log, go to the details page for a resource, click Data Collection , and select Actions > Open Logs in the Probe section on the Data Collection page.
Status*	One of the following statuses is detected on a fabric:
	Not Normal An error or warning status was detected on the fabric or its internal resources.
	Warning A warning status was detected on the fabric or its internal resources.
	Error (default) An error status was detected on the fabric or its internal resources.
	Unreachable One or more of the monitored resources for a fabric are not responding.

Triggering conditions for server alerts

You can set up IBM Spectrum Control so that it examines the attributes and capacity of a server and notifies you when changes are detected.

Alerts can notify you of general changes and capacity changes on the following resources:

- "Servers" on page 220
- "Controllers" on page 222

- "Disks" on page 223
- "Disk Groups" on page 223
- "File Systems and Logical Volumes" on page 224
- "Paths" on page 225
- "Shares" on page 225
- "Triggering conditions for the IBM Spectrum Control server" on page 226

Important: A number of other attributes are available for alerts and are based on the key properties of a server. To view a complete list of server attributes upon which you can alert, go to Settings > Alert Policies. Double-click a default policy for a server. Click Edit Alert Definitions on the Alert Definitions tab. View the attributes that are available in the general and capacity categories. Note that the attributes that are automatically configured for alerts in the default alert policies, or default alerts, have a status of Active.

In the tables, default alerts are marked with an asterisk (*).

Tip: For capacity attributes, you can generate alerts when the amount of storage is greater than, less than, or equal to a specified value. You can also determine the unit of measurement for the attribute, such as KiB, MiB, GiB, or TiB.

Servers

Table 88. Triggering attributes and conditions for general changes on servers	
General Attributes	Descriptions of Triggering Conditions
Agent State*	A Storage Resource agent is in one of the following states:
	Not Normal An error or warning state was detected on a Storage Resource agent.
	Warning A warning state was detected on a Storage Resource agent. For example, a warning state might occur when an agent must be upgraded to the same version level as the IBM Spectrum Control server to which it is communicating.
	Error (default) An error state was detected on a Storage Resource agent. For example, an error state might occur when an agent was not able to be upgraded.
Last Successful Probe	A specified amount of time has passed since a probe was able to collect data about a server. You can use this alert to be notified when up-to-date configuration and status data is not being collected about a server and its existing data might be stale. Data collection might be interrupted or not occur if the resource, network, or IBM Spectrum Control server are unavailable.

General Attributes	Descriptions of Triggering Conditions
Probe Status*	One of the following statuses is detected for a probe:
	Not Successful An error or warning occurred during data collection. This status indicates that a probe did not collect any data, or only collected a partial set of data about a resource.
	Warning A probe completed, but might not have collected a complete set of data. This status might occur if data cannot be collected about one or more of the internal resources of a resource.
	Error (default) A probe did not complete when it attempted to collect asset data about the resource. This status might occur if the resource cannot be reached during data collection.
	For details about why a specific status occurred, check the log for the probe. To check the log, go to the details page for a resource, click Data Collection , and select Actions > Open Logs in the Probe section on the Data Collection page.
Status*	One of the following statuses is detected on a server:
	Not Normal An error or warning status was detected on the server or its internal resources.
	Warning A warning status was detected on the server or its internal resources. For example, a warning status might occur when an HBA or HBA to a server node is newly discovered, is missing, or is rediscovered.
	Error (default) An error status was detected on the server or its internal resources. For example, an error status might occur when a server goes offline, or a server disk is disconnected or partially disconnected if it has multiple paths and one of the paths is disconnected.
	Unreachable One or more of the monitored resources for a server are not responding. This status might be caused by a problem in the network or by a Storage Resource agent that is no longer running and did not communicate that it was shutting down.

Table 89. Triggering attributes and conditions for capacity changes on servers	
Capacity Attributes	Descriptions of Triggering Conditions
Available Drive Capacity (Previously known as Available Disk Space)	The unused disk capacity on the local and SAN-attached storage for the server. SAN-attached storage is assigned to the server from storage systems.
Available File System Capacity	The amount of unused capacity in the file systems on the server. Available file system capacity does not include capacity that is reserved for the operating system. For example, the available capacity for tmpfs on UNIX operating systems is not included in this value.

Table 89. Triggering attributes and conditions for capacity changes on servers (continued)	
Capacity Attributes	Descriptions of Triggering Conditions
Drive Capacity (Previously known as Total Disk Space)	The total disk capacity for all the local and SAN-attached storage on the server. SAN-attached storage is assigned to the server from storage systems.
File System Capacity	The amount of file system capacity on the server.
File System Capacity from Storage Systems	The amount of file system capacity that is assigned to the server from storage systems.
	The file system capacity from storage systems is only available when SAN-attached storage is assigned to the server.
Mapped SAN Capacity (Previously known as Assigned SAN Space)	The amount of disk capacity that is assigned to the server from storage systems.
	The disk capacity from storage systems is only available when SAN-attached storage is assigned to the server.
Used Capacity	The amount of used disk capacity on the local and SAN-attached storage for the server. SAN-attached storage is assigned to the server from storage systems.

Controllers

Table 90. Triggering attributes and conditions for changes on disk controllers	
Controller Attributes	Descriptions of Triggering Conditions
Driver VersionFirmwareROM Version	The version of the driver, firmware, or read-only memory (ROM) on a disk controller changes. You can use a number of operators to determine when you are notified of a version change, such as when the driver, firmware, or ROM is, or is not, a specific version, or when the version number contains a specific value. Use this alert for HBAs only.
Last Data Collection	A specified amount of time since data was collected about a controller. Use this alert to be notified if data is not being collected about a controller or if the existing data is becoming too stale.
New Disk Controller	A disk controller is detected for the first time. Use this alert to be notified of hardware additions on servers.
Removed Disk Controller	A previously monitored disk controller can no longer be found. Historical data about the controller is retained, but no current data is being collected. Use this alert to be notified if a controller is removed or becomes unavailable.
Status*	One of the following statuses is detected on a disk controller: Not Normal An error or warning status was detected on the controller. Warning A warning status was detected on the controller. Error (default) An error status was detected on the controller.

Disks

Table 91. Triggering attributes and conditions for general changes on disks	
General Attributes	Descriptions of Triggering Conditions
Firmware	The version of the Licensed Internal Code on the disk changes. You can use a number of operators to determine when you are notified of a firmware change, such as when the firmware is, or is not, a specific version, or when the version number contains a specific value.
Multipathing Policy	The multipathing policy that is in effect for a disk. For example, you can be notified when the policy changes, or when the policy is Round Robin, Load Balancing, Failover Only, or other another policy.
New Disk	A disk is detected for the first time. Use this alert to be notified of hardware changes on servers or hypervisors.
Paths	The number of access paths that are associated with the disk falls outside a specified range, or is equal to or not equal to a specified value.
Removed Disk	A previously monitored disk can no longer be found. Historical data about the disk is retained, but no current data is being collected. Use this alert to be notified if a disk is removed or becomes unavailable.
Status*	One of the following statuses is detected on a disk: Not Normal An error or warning status was detected on the disk. Warning A warning status was detected on the disk. Error (default) An error status was detected on the disk.

Table 92. Triggering attributes and conditions for capacity changes on disks	
Capacity Attributes	Descriptions of Triggering Conditions
Available Drive Capacity (Previously known as Available Disk Space)	The unused capacity on a disk that is attached to the server.
Capacity	The total amount of storage capacity on a disk that is attached to the server.
Used Capacity	The amount of used storage capacity on a disk that is attached to the server.

Disk Groups

Table 93. Triggering attributes and conditions for general changes on disk groups (volume groups)	
General Attributes	Descriptions of Triggering Conditions
Deleted Volume Group	A previously monitored volume group can no longer be found. Historical data about the volume group is retained, but no current data is being collected. Use this alert to be notified if a volume group is deleted or becomes unavailable.
New Volume Group	A volume group is detected for the first time.

Table 93. Triggering attributes and conditions for general changes on disk groups (volume groups) (continued)	
General Attributes	Descriptions of Triggering Conditions
Status*	One of the following statuses is detected on a disk group:
	Not Normal An error or warning status was detected on the disk group.
	Warning A warning status was detected on the disk group.
	Error (default) An error status was detected on the disk group.

Table 94. Triggering attributes and conditions for capacity changes on disk groups (volume groups)	
Capacity Attributes	Descriptions of Triggering Conditions
Available Capacity	The unused storage capacity on a server disk group.
Used Capacity	The amount of used storage capacity on a server disk group.
Volume Group Capacity	The total amount of storage capacity on a server volume group. This value is inclusive of all storage capacity and applies to all capacity values related to volume groups.

File Systems and Logical Volumes

Table 95. Triggering attributes and conditions for general changes on file systems and logical volumes	
General Attributes	Descriptions of Triggering Conditions
Available Inodes	The number of unused inodes on file systems on the operating system changes, falls outside a specified range, or is equal to or not equal to a specified value.
Deleted File System	A previously monitored file system is deleted or unmounted from a server. Historical data about the file system is retained, but no current data is being collected. This attribute applies to file systems on the following resources:
	Storage systems that are configured for file storage, including Storwize V7000 Unified
	Servers that are managed by Storage Resource agents
Deleted Logical Volume	A previously monitored logical volume can no longer be found. Historical data about the logical volume is retained, but no current data is being collected. Use this alert to be notified if a logical volume is removed or becomes unavailable.
New File System	A file system was detected for the first time. This alert applies to file systems on the following resources:
	Storage systems that are configured for file storage, including Storwize V7000 Unified
	Servers that are managed by Storage Resource agents
New Logical Volume	A logical volume is detected for the first time.

Table 95. Triggering attributes and conditions for general changes on file systems and logical volumes (continued)	
General Attributes Descriptions of Triggering Conditions	
Used Inodes	The number of used inodes on file systems on the operating system changes. You can use a number of operators to determine when you are notified, such when the number of used inodes falls outside a specified range, or is equal to or not equal to a specified value.

Table 96. Triggering attributes and conditions for capacity changes on file systems and logical volumes		
Capacity Attributes	Descriptions of Triggering Conditions	
Available File System Capacity	The amount of unused storage capacity on a file system on the server disk.	
File System Capacity	The total amount of storage capacity on a file system on the server disk.	
Logical Volume Capacity	The total amount of storage capacity on a logical volume on the server disk.	
Used File System Capacity	The capacity on a file system on the server disk.	
Used Capacity	The percentage of used storage capacity on a file system or logical volume on the server disk.	

Paths

Table 97. Triggering attributes and conditions for general changes on paths		
Path Attributes	Descriptions of Triggering Conditions	
Deleted Path	A previously monitored access path for a server disk can no longer be found. This change might or might not affect the availability of the disk because there might be more than one path available.	
New Path	An access path for a disk is detected for the first time.	
Status*	One of the following statuses is detected on a path: Not Normal An error or warning status was detected on the path. Warning A warning status was detected on the path. Error (default) An error status was detected on the path.	

Shares

Table 98. Triggering attributes and conditions for general changes on shares		
Share Attributes	Descriptions of Triggering Conditions	
Deleted Share	A previously monitored share can no longer be found. Historical data about the share is retained, but no current data is being collected. Use this alert to be notified if a share is removed or becomes unavailable.	
New Share	A new share was detected for the first time.	

Triggering conditions for the IBM Spectrum Control server

The server on which IBM Spectrum Control is installed is automatically monitored for conditions that might cause an interruption in product functions. When these conditions are detected, alerts are triggered and shown on the **Home** > **Alerts** page. You do not need to manually define alerts for these product-related conditions; they are automatically enabled.

Table 99. Triggering conditions for the IBM Spectrum Control server		
Triggering Condition	Explanation	Related Error Message
Database unavailable	The product database is not available. This database is the repository for information that is collected about the monitored resources in your environment.	ALR4112E, ALR4113E
High memory usage*	A high amount of memory is being used by a server process and might cause stability problems.	ALR4103W
Database alarm*	The system database or the database manager that hosts the product's database repository is reporting an alarm.	ALR4104W
High workload	The workload queue for the Device server is high and might cause performance issues.	ALR4105W
High number of external events	The server is receiving a high number of external events, such as CIM indications or SNMP traps. The high number of events might cause performance issues.	ALR4106W

Alert notifications and actions

Specify how you are notified when alert conditions are detected on resources, and define actions to take as a result of those alerts. These settings are defined globally for all resources, and can be overridden for a specific alert definition, for all alert definitions that apply to a specific resource, or for an alert policy.

Notifications and override behavior

You can specify notification settings at the following levels:

- Globally, for all alerts.
- In alert policies, for groups of resources of the same type.
- In individual alert definitions, which can be part of an alert policy, or can be specific to a resource.

Notification settings at the lower levels override the settings at the higher levels. For example, if you want to use email for alert notifications, you can specify email notification settings like this:

- 1. If you want all email notifications to go to a specific set of email addresses, specify email addresses for global alert notifications.
- 2. If you're using alert policies, you can also specify a list of email addresses to notify that are applicable to that policy. These addresses override the global setting.
- 3. For an alert definition, you can specify a set of email addresses to notify. This email addresses override the alert policy and global settings.

Scenarios for the different levels of alert notification

A small organization might specify only global email addresses if they want to send all alert notifications to one team.

A larger organization can have multiple storage environments that are managed by different teams. They can create alert policies for each environment and specify different email addresses for each alert policy. For example, they might create alert policies for their mainframe, VMware, and AIX environments, and

specify different email addresses for each of those policies. These email addresses override any global email addresses.

If you want to notify a specific person when a specific condition occurs, you set an email address in an individual alert definition. This email address overrides any alert policy or global email addresses. For example, an operations team might review all of the alerts, but might want to notify their storage architect if capacity in a pool exceeds 80%.


Specifying the frequency of alert notifications

You can use the Notification Frequency settings to avoid triggering too many alerts for some conditions. You can select one of the following options:

Option	Description	
Send every time condition is violated	Receive alert notifications whenever an alert condition is violated.	
Send once until problem clears	Receive one notification for a violation, even if the condition is violated multiple times.	
Send every time_period	Receive one notification when an alert condition is initially violated. The alert is suppressed and no notifications are sent until both the specified time has passed and the alert is triggered again.	

You can also select to only send notifications if the violation is not cleared for longer than a time period that you specify. Select **Send after condition is violated for** *time_period*, then specify the time period.

For example, you might want to be notified about an alert only after the condition has been violated for 20 minutes, and you only want to be notified about the alert once until the problem clears. In this case, you can set up an alert with the following notification frequency settings:



Restrictions:

- For some attributes, not all the notification frequency options are available. Specifically, you cannot change the notification frequency for the following attributes: New *resource*, Removed or Deleted *resource*, Last Successful Probe, and Last Successful Monitor.
- For attributes that use the **changes** operator (for example, the Firmware attribute), only the **Send every time condition is violated** and **Send every time_period** notification frequency options are available.

Specifying alert actions and notification settings in the GUI

Table 100. Locations where you can specify alert actions and notification settings		
Task	Location in GUI	
Specify the global notification settings for all alerts	Global Alert Notifications To specify the global notification settings, go to Configuration > Notification Settings, and specify the details for email, SNMP, and Tivoli Netcool/OMNIbus. These notification settings are applied globally to all alert definitions, unless overridden.	
	Important: You must specify the global alert notifications settings for each type of alert notification that you want to enable.	
Specify the notification settings for an alert policy	Alert policies page To specify the notification settings for an alert policy, go to Settings > Alert Policies. Double-click the policy whose notification settings you want to specify. Click Edit Policy Notifications.	
	Specify the details for email, SNMP, and Tivoli Netcool/OMNIbus. These notification settings are applied to all of the alert definitions for all resources in the policy, unless overridden.	
Specify the notification settings for all the alert definitions that apply to a resource that is not managed by an alert policy	Resource details page For example, to specify the notification settings for a block storage system, go to Storage > Block Storage Systems. Double-click the resource for which you want to define alerts. In the General section of the resource details page, click Alerts. Then click Edit Notifications.	
	These notification settings are applied to all the alert definitions that are specified for the selected storage system. You can override the global settings by specifying different settings for each alert definition.	
Specify alert actions and notification settings for a specific alert definition	Alert definitions page, Resource details page For example, to specify the notification settings for an alert definition in an alert policy, go to Settings > Alert Policies. Double-click the policy whose notification settings you want to specify. Click Edit Alert Definitions.	
	Alternatively, to specify the notification settings for an alert definition for a resource that is not managed by an alert policy, go to Storage > Block Storage Systems . Double-click the resource for which you want to define alerts. In the General section of the resource details page, click Alerts . Then click Edit Alert Definitions .	
	Click on the alert definition to show advanced options. Specify the details for email, SNMP, and Tivoli Netcool/OMNIbus. The notification settings that you define override any global notification settings, any policy settings, and any settings for the resource.	

Alert actions

You can specify that the following actions are taken when alert conditions are detected on monitored resources:

Run script

Run a script when an alert is triggered for the condition. Use a script to call external programs or run commands that take action as the result of an alert. By using a script, you can automatically address potential storage issues when they are detected to avoid unplanned downtime or performance bottlenecks. Learn more.

Email

Send emails to specific email addresses when an alert is detected on a monitored resource.

Netcool or OMNIbus EIF Probe Server

Send alert notifications to a Netcool server or OMNIbus EIF probe server within your environment that was configured to receive IBM Spectrum Control alerts.

SNMP trap

Generate SNMP trap messages to any network management station (NMS), console, or terminal when an alert condition is detected. System administrators must set up their SNMP trap ringer with the provided management information base (MIB) files to receive SNMP traps from the product.

Windows event log or UNIX syslog

Write alert messages to the OS log. If you already have an administrator monitoring OS logs, this method is a way to centralize your priority messages for quick notification and viewing.

Related tasks

Defining alerts

Define how and when you are alerted to changes in your storage environment. You can define alerts for alert policies, for individual resources, or for sets of resources that are included in an application or general group.

How scripts are run

After an alert is triggered, scripts can be run from the server where IBM Spectrum Control is installed or from any server where a Storage Resource agent is deployed. For some alert conditions, scripts can be specified to run on the server where the change was detected.

You must include the script files that you want to run in the /scripts directory on the server where IBM Spectrum Control is installed. The default path of /scripts is determined by the operating system of the server:

- Windows: C:\Program Files\IBM\TPC\Data\scripts
- UNIX or Linux: /opt/IBM/TPC/Data/scripts or /usr/IBM/TPC/Data/scripts

Tip: When you run a script against a NAS filer, the script is run from the Storage Resource agent that is assigned to the file system where the condition was detected.

Script parameters

You can run scripts in response to alerts. Parameters for the conditions and attributes in an alert definition are included in the script. Each parameter is associated with a number, such as \$1, \$2, \$3, and so on. These parameters are used when the alert is triggered.

The following examples show the parameters that are included in a script for different alerting scenarios.

Alerting example: Storage System Pool Used Capacity >= 10

```
$1 = Pool
$2 = tpcvm4-3
$3 = "Storage System"
$4 = SVC-storage1
$5 = "Storage System Pool Used Capacity Percent >= 10 %"
$6 = "75.97 %"
```

Alerting example: Storage System Write Data Rate <= 1,100 MiB/s

```
$1 = ""

$2 = ""

$3 = "Storage System"

$4 = SVC-storage1
```

```
$5 = "Write Data Rate <= 1,100 MiB/s"
$6 = "0 MiB/s"
```

Alerting example: Pool Total I/O Rate <= 445 ops/s

```
$1 = Pool
$2 = scv5k2c0_pool
$3 = "Storage System"
$4 = SVC-storage1
$5 = "Total I/O Rate - overall <= 445 ops/s"</pre>
```

Alerting example: Custom alert with the multiple conditions

Conditions:

- Storage System Volume Capacity >= 1,000,000 KiB
- Storage System Volume Compressed is yes
- Storage Volume Compression Savings Percent >= 13

Parameters:

```
$1 = Volume
$2 = vol1
$3 = "Storage System"
$4 = SVC-storage1
$5 = "Storage System Volume Capacity >= 1,000,000 KiB"
$6 = "2,097,152 KiB"
$7 = Volume
$8 = vol1
$9 = "Storage System"
$10 = SVC-storage1
$11 = "Storage System Volume Compressed is yes"
$12 = yes
$13 = Volume
$14 = vol1
$15 = "Storage System"
$16 = SVC-storage1
$17 = "Storage Volume Compression Savings Percent >= 13 %"
$18 = "17.25 %"
```

Restriction: For some alerts, such as custom alerts with multiple conditions and attributes, the number of parameters in the associated script might exceed nine. In Windows, only the first nine parameters in a script can be referenced by number (parameter %1 to parameter %9). Additional parameters are ignored.

How script names are resolved

When you run scripts in response to alerts, the names of those scripts are resolved by using the following criteria:

- The operating system of the server where the product is installed.
- The script name that you specify on the **Script name** field when you define or edit an alert definition.

The operating system of the server or Storage Resource agent where you choose to run a script determines how that script that is run:

UNIX, Linux operating systems

A Storage Resource agent that is deployed on a UNIX or Linux operating system does not run scripts that have an extension. 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 searches for a file with the name backup. If the script exists, the first line of the script is read and the appropriate interpreter is used to run the script. If a file cannot be found, no action is taken.

Windows operating systems

A Storage Resource agent that is deployed on a Windows operating system runs scripts that have an extension. The extension of a script file determines which interpreter is used to run the script. If you specify a script name that does not contain an extension, the agent searches for the file name in the following order (based on its extension): .com, .exe, .bat, .cmd, .vbs, .vbe, .js, .jse. If a file cannot be found, no action is taken.

Tip: You can 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 in the /scripts directory. For example, if you want to run a provisioning script against Storage Resource agents on UNIX or Linux and Windows servers, you must have two versions of that file (provision.bat and provision) in /scripts.

Viewing information about resources

You can view detailed information about the resources that are monitored by IBM Spectrum Control. Resources include storage systems, volumes, pools, servers, hypervisors, fabrics, and switches. You can also view information about internal and related resources.

Remember: Data collection jobs retrieve information about monitored resources. To ensure that the most up-to-date information about resources is available in the GUI, run data collection jobs regularly. To learn about how to collect data, see "Collecting data" on page 138.

How information is organized

Information about monitored resources is organized into different pages in the GUI. You can use the pages to view summary information about all types of monitored resources, such as storage systems, fabrics, and switches. You can see what resources are consuming storage and the alerts that are being detected on those resources. You can also view a chart of the most active storage systems in your environment.

Dashboard view

The dashboard is view displayed when you start a new session of IBM Spectrum Control or when you go to **Home** > **Dashboard** in the menu bar. Use the dashboard to view the following summary information about a storage environment

- Condition of storage: block, file, and object storage systems
- Condition of storage consumers: hypervisors and servers
- Condition of network resources: fabrics and switches
- Number error, warning, and information alerts that are generated on resources
- Most active storage systems in your environment.

You can click on resource icons in the dashboard to add more resources for monitoring and to see detailed information about resources that IBM Spectrum Control is currently monitoring. Position the mouse pointer over the lines in the Most Active Storage Systems chart to see performance information about specific resources. For more information about the dashboard, see "Dashboard view" on page 233.

Resource list pages

Use a resource list page as a central location for viewing information about the top-level resources of a specific type. For example, use the **File Storage Systems** page to view a list of the file storage systems that are monitored by IBM Spectrum Control. Each resource list page corresponds to a specific type of resource, which includes block storage systems, file storage systems, object storage systems, volumes, pools, servers, hypervisors, switches, and fabrics. You can view the following information on resource list pages:

- Overall condition of the resources
- Key information about the resources
- · Alerts that were generated for the resources
- · Jobs that were run for the resources

To learn about how to access the list pages for resources, see <u>"Viewing information about top-level</u> resources" on page 236.

Note: Volumes and pools are internal resources of storage systems, but have their own resource list pages so that you can view them all in a central location.

Resource detail pages

Use resource details pages to view information about resources that are internal or related to a top-level resource. For example, use the File Storage System details page to view detailed information about a specific file storage system. You can view the following information on resource details pages:

- The overall capacity and storage usage for a resource
- · The status of resources
- · Key information about a resource
- · Alerts that were generated for a resource
- Data collection that was run for a resource
- Information about components that exist in top-level resources
- Information about the resources that are related to another resource

To learn about how to access detail pages for resources, see "Viewing information about internal, object, and related resources" on page 237.

Alerts page

Use the **Alerts** page to view and manage the alert conditions that were detected on monitored resources. Specifically, you can complete the following actions:

- View the overall number of alerts that were detected on monitored resources. This information is presented in a bar chart for the following types of alerts:
 - Database
 - Fabric
 - Hypervisor
 - Other
 - Performance
 - Server
 - Storage system
 - IBM Spectrum Control
- View information about specific alerts, including severity, last occurrence, and the resources where they
 were detected.
- Acknowledge alerts that were reviewed but are not yet resolved.
- Unacknowledge alerts.
- · Remove alerts from the page.

To access the **Alerts** page, go to **Home > Alerts** in the menu bar.

Performance Monitors page

IBM Spectrum Control can collect information about the performance of monitored storage systems and switches. This information includes key performance metrics. On the **Performance Monitors** page, you can view detailed information about performance monitors that are collecting data. Use this information to identify performance monitors that are encountering problems when they run. You can also view the metrics that are collected by a performance monitor for a resource.

To access the **Performance Monitors** page, select **Home > Performance Monitors**.

Tasks page

Use the **Tasks** page to view and manage the tasks that IBM Spectrum Control uses to optimize resources and provision storage. Tasks are created when you complete any of the following actions:

- Analyzing tiering by moving volumes to higher or lower tiers.
- Transforming storage by moving or converting volumes.
- Balancing pools by distributing the workload of volumes across pools on the same tier.

To access the **Tasks** page, go to **Home > Tasks**.

System Management page

Use the **System Management** page to view information about the overall condition of IBM Spectrum Control. You can view information about the servers on which IBM Spectrum Control is installed. This information includes component server and database status, certain server alerts and database connection alerts, server file-system capacity information, and remote volume-performance information.

To access the **System Management** page, go to **Home > System Management** in the menu bar.

Reporting

Use the optional Cognos Analytics reporting tool in the GUI to view detailed reports about resources. The following reports are available:

Predefined reports

These reports are included with IBM Spectrum Control. The following types of predefined reports are available:

- Capacity and relationships of resources
- Performance
- · Historical space
- Storage tiering

Custom reports

Unlike predefined reports, you can select the information that you want to include in custom reports. Custom reports can contain detailed information about the relationships between monitored resources, the properties of monitored resources, and detailed information about the performance of monitored resources.

To learn about how to use the Cognos Analytics reporting tool, see Peporting with Cognos Analytics.

Dashboard view

Use the dashboard to monitor the overall condition of monitored resources and identify potential problem areas in a storage environment. The dashboard also displays alerts, as well as a chart showing the most active storage systems.

Viewing the status of monitored resources

The dashboard provides status information for the resources that IBM Spectrum Control is monitoring.

The overall status for each resource type

A status symbol and number are shown below each resource icon on the dashboard. This symbol represents the most critical status that was detected on the internal resources for a resource type. For example, if 20 storage systems are monitored, and an error was detected on a volume for one of

those storage systems, a 1 status is shown below the storage systems icon or label in the dashboard. If no errors, warnings, or unreachable statuses were detected on the monitored resources, then a green symbol is shown without a number.

Use the status to quickly determine the condition of your monitored resources and if any actions must be taken. On the dashboard, the following statuses might be shown for a resource type:



An error status was detected on one or more of the monitored resources for a resource type. Error statuses represent serious problems on a resource. Resolve these problems as soon as possible. View the details of a resource to learn more about the status of its internal resources.



Unreachable (orange)

One or more of the monitored resources for a resource type are not responding. This status might be caused by a problem in the network or by a Storage Resource agent that is no longer running and did not communicate that it was shutting down.



Warning (yellow)

A warning status was detected on one or more of the monitored resources for a resource type. Warning conditions are not critical, but represent potential problems. View the details of a resource to learn more about the status of its internal resources.



Normal (green)

No warnings or errors were detected on the monitored resources for a resource type.



Unknown status (gray)

One or more of the monitored resources for a resource type have an unknown status. This status might occur if no data was collected about a resource. To change an unknown status, use a data collection job to collect status information about the resource.

For information about how status is determined, see "How the condition of a resource is determined" on page 248.

The number of statuses that were detected and acknowledged for each resource type

To view the total number of statuses for a resource type, move the mouse pointer over the related status symbol. Hover help shows the total number of occurrences for each status, and the number of statuses that were acknowledged by a IBM Spectrum Control user. An acknowledged status indicates that a status was reviewed and is either resolved or can be ignored.

Viewing the status of monitored storage systems

Storage systems can be configured for block storage, file storage, object storage, a combination of block and file storage, or a combination of file and object storage. The top left section of the dashboard shows the type, number, and status of the storage systems that IBM Spectrum Control is monitoring. The number in the circle next to the storage system type shows how many storage systems of that type are being monitored. The status icon below each storage system type shows the status of the monitored systems. Click the storage system type, for example Block Storage Systems, to see more detailed status and information about those storage systems. Use the information in the horizontal bar charts to view details about the used and free space of the monitored storage systems in your environment:

Storage systems that are configured for block data

The following information is shown for monitored storage systems that are configured for block data:

Used Block Storage Space

Amount of space that is being used for block data on all the storage systems that are configured for block data.

Available Block Storage Space

Amount of unused space on all the storage systems that are configured for block data.

Storage systems that are configured for file data

The following information is shown for monitored storage systems that are configured for file data:

Used File System Capacity

Amount of file system capacity that is being used by files and directories on the storage systems and filers that are configured for file data.

Available File System Capacity

Amount of unused storage space on the storage systems and filers that are configured for file data.

Storage systems that are configured for object data

The following information is shown for monitored storage systems that are configured for object data:

Used Object Storage Space

Amount of storage space that is used on all the object storage systems.

Available Object Storage Space

Amount of storage space that is available on all the object storage systems.

Viewing the status of monitored resources that are consuming storage

The diagram in the top middle section of the dashboard provides a visual representation of the monitored resources that are consuming storage. You can click the icons in this section to add resources of that type or to see detailed information about the resources you have already added. The number next to each icon shows the number of those resources that are being monitored by IBM Spectrum Control. If IBM Spectrum Control is monitoring resources of a particular type, such as servers, you will see a status icon that shows the overall status of those types of resources.

The types of resources that are consuming storage

Each type of resource in the diagram is represented by an icon. To learn more detailed information about the individual resources that are monitored, click the icon in the diagram. For example, if you click the departments icon, a new page shows detailed information about the individual departments that are being monitored.

The number of resources that are monitored for each resource type

A number in a circle is shown next to each resource icon. This number represents the number of monitored resources for each resource type. For example, if 4 file storage systems are monitored, the value 4 is shown. This number excludes resources that are not monitored by IBM Spectrum Control.

Tip: If no resources are being monitored for a resource type, click its icon on the dashboard to add a resource of that type. For example, if the dashboard shows a 0 next to the departments icon, you can click the icon or the circle for departments to add a department. The GUI guides you though the process for adding the resource for monitoring. For more information about adding resources, see "Adding resources" on page 115.

Viewing the status of monitored network fabrics and switches

The top right section of the dashboard shows the number and status of the fabrics and switches that IBM Spectrum Control is monitoring. The number in the circle next to the fabric or switch icon shows how many fabrics and switches are being monitored. The status icon below the fabric or switch shows the status of those entities. Click the icon to see more detailed status and information about those fabrics and switches. The dashboard shows the following information for the monitored fabrics and switches in your environment:

Fabrics

- The number of fabrics that are being monitored.
- The most critical status of the monitored fabrics. If more than one fabric is monitored, the dashboard shows the most critical status level of those fabrics.

Switches

- The number of switches that are being monitored.
- The most critical status of the monitored switches. If more than one fabric is monitored, the dashboard shows the most critical status level of those switches.

Viewing the status of alerts that were detected on monitored resources

Alerts are generated when IBM Spectrum Control detects certain conditions or events on monitored resources. The Alerts section of the dashboard shows a status summary of the alerts that were detected in your environment during a specified time period. This summary includes the number of alerts that have the following statuses: Critical, Warning, and Informational.

Use the following actions to manage the display of statuses for alerts:

- To select the time period in which you want to view alert statuses, click the down arrow and select a time period. You can view the status of all alerts that were detected in the last hour, last day, or last week.
- To view more detailed information about alerts, click **View all alerts** to access the **Alerts** page.

Viewing the most active storage systems in your environment

The Performance chart provides information about the most active storage systems in your storage environment. Position your cursor over a data point in the chart to see the total I/O rate for a particular storage system. The I/O rate is displayed in operations per second (ops/s) on the left vertical axis, and in milliseconds per operation (ms/ops) on the right vertical axis. The horizontal axis shows the time of the performance data in hourly increments.

Viewing information about top-level resources

Use a resource list page as a central location for viewing information about all the monitored resources of a specific type. Resources that have a list page include storage systems, volumes, pools, shares, servers, hypervisors, switches, and fabrics.

1. In the menu bar, hover the mouse pointer over the menu item for the type of resource that you want to view.

You can access the following types of resources:

- Storage Block storage systems, file storage systems, Object Storage systems, volumes, pools, shares
- **Servers** Servers, hypervisors
- Network Switches, fabrics
- 2. Select the type of resource that you want to view.

For example, if you want to view information about servers, go to **Servers** > **Servers**.

A list of monitored servers is displayed on Servers page.

- 3. Locate the resource that you want to view in the list.
 - Each column in the list shows information about the resource.
- 4. Optional: To view key attributes of a specific top-level resource without leaving the page, right-click the resource and select **View Properties**.
 - A properties notebook is displayed. Information in a properties notebook can include storage statistics, hardware attributes, volume assignments, disk assignment, and performance metrics.
- 5. Optional: To view more detailed information about a specific top-level resource, including its internal and related resources, right-click the resource and select **View Details**.
 - A resource details page is displayed. The page includes information about storage statistics, alerts, jobs, internal resources, and related resources.

A resource list page shows information about the monitored resources for the selected type. You can view the following information on the page:

- The overall status of the resources for a specific type.
- · Key information about the resources.
- · Alerts that were generated for the resources.
- · Jobs that were run for the resources.

Tip: To view descriptions of the information that is available on a resource list page, click the **Help** icon in the upper right of the page.

To view information about rollup resources, you must be in rollup mode. To enter rollup mode, move the

mouse pointer over the rollup icon in the menu bar and select **Enter rollup mode**.

Related reference

"Resources that you can monitor" on page 111

You can use IBM Spectrum Control to monitor storage systems, servers, hypervisors, fabrics, and switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

Viewing information about internal, object, and related resources

Use resource details pages to view information about a specific top-level resource and its internal, object, and related resources.

1. In the menu bar, hover the mouse pointer over the menu item for the type of resource that you want to view.

You can access the following types of resources:

- Storage: block storage systems, file storage systems, object storage systems, volumes, pools, shares
- **Servers**: servers, hypervisors
- Network: switches, fabrics
- 2. Select the type of the resource that you want to view.

A resource list page shows the monitored resources for the selected type. For example, if you want to view a block storage system, go to **Storage** > **Block Storage Systems**. A list of monitored block storage systems is displayed.

- 3. In the list of resources, right-click the resource that you want to view and select **View Details**.

 A resource details page is displayed. This page includes information about storage statistics, alerts, jobs, internal resources, object resources, and related resources.
- 4. To view information about the internal resources of the top-level resource, click the name of the internal resource in the **Internal Resources** section.
 - A list of the internal resources is displayed. Each column in the list shows information about the resources. For example, on the details page for a server, click **Controllers** to view a list of the controllers that are associated with the server.
- 5. Optional: To view key attributes of a specific internal resource without leaving the page, right-click the resource and select **View Properties**.
- 6. To view information about the object resources in an object storage system, click the name of the object resource in the **Object Resources** section.
 - Object resources are associated only with object storage systems, such as IBM Spectrum Scale that is configured for object storage.
- 7. To view information about resources that are related to the top-level resource, click the name of the related resource in the **Related Resources** section.
 - A list of the related resources is displayed. For example, on the server details page, click **Block Storage Systems** to view a list of the block storage systems that have at least one disk (volume) assigned to the server.
- 8. Optional: To view key attributes of a specific related resource without leaving the page, right-click the resource and select **View Properties**.
- 9. Optional: To view the resource details page for a related resource, right-click the resource and select **View Details**.

A resource details page shows information about a top-level resource. The information is available in the following links and sections on the page:

- Overview: The overall capacity and storage usage for a resource.
- Properties: Key information about a resource.
- Alerts: Alerts that were generated for a resource.
- Tasks: Jobs that were run for a resource.
- Internal Resources: Information about the components that exist in a higher-level resource.
- **Object Resources**: Information about the resources in a IBM Spectrum Scale cluster, such as accounts and containers, that enable the IBM Spectrum Scale cluster to be used as an object storage system.
- **Related Resources**: Information about the resources that are related to another resource. A related resource is external to a higher-level resource, but is related to it through assigned storage, a network connection, or virtual hosting.

Tip: To view descriptions of the information that is available on a resource details page, click the **Help** icon **1** in the upper-right corner of the page.

Related reference

"Resources that you can monitor" on page 111

You can use IBM Spectrum Control to monitor storage systems, servers, hypervisors, fabrics, and switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

Overview charts

Monitor the activity and storage usage of your storage systems. Use the charts and key metrics to gauge the workload activity of your storage resources and to check whether you have sufficient capacity.

Review the charts and key information about the activity and availability of capacity:

- For block storage systems
- For file storage systems
- For object storage systems

Block storage systems

To monitor storage usage, detect capacity shortages, and plan your capacity needs, use the charts on the **Overview** page.



Figure 3. Capacity and activity charts

Three types of charts are provided:

- · Capacity overview charts
- · Capacity by storage resource charts
- · Workload activity charts

The **Capacity** chart at the top of the **Overview** page shows how much capacity is used and how much capacity is available for storing data.

The **Provisioned Capacity** chart shows the written capacity values in relation to the total provisioned capacity values before data reduction techniques are applied. The following values are shown:

- The capacity of the data that is written to the volumes as a percentage of the total provisioned capacity of the volumes.
- The amount of capacity that is still available for writing data to the thin-provisioned volumes in relation to the total provisioned capacity of the volumes. Available capacity is the difference between the provisioned capacity and the written capacity, which is the thin-provisioning savings.

A breakdown of the total capacity savings that are achieved when the written capacity is stored on the thin-provisioned volumes is also provided.

In the capacity overview chart, a horizontal bar is shown when a capacity limit is set for the storage system. Just hover over the chart to find out what the capacity limit is and how much capacity is left before the capacity limit is reached.

To get a breakdown of the capacity usage by pool or volume, click the links.

Try it: Click **View capacity by pool**. You get a capacity breakdown of the storage system by pool, and you can see the recent fill and growth rates for each pool in the storage system.

Depending on the type of storage system, the following activity and capacity by storage resource charts are shown:

Table 101. Capacity by storage resource	
Chart Name	Purpose
Capacity	Use this chart to monitor the available, used, and provisioned capacity in your storage system over the last 30 days.
Capacity by Host	Use this chart to monitor the current capacity of the storage system that is mapped to hosts.
Capacity by Pool	Use this chart to monitor the current available and used capacity of your largest pools.
Capacity by Volume	Use this chart to monitor the current used and unused provisioned capacity of your largest volumes.
Capacity by Tier	Use this chart to monitor the current used and available capacity of the pools that are assigned to tiers.

Table 102. Activity charts	
Chart Name	Purpose
Overall system activity	Use this chart to monitor the overall activity of your storage system, which is broke down into total I/O rate, read I/O rate, and write I/O rate in operations per second over the last 24 hours.
Most Active Nodes	Use this chart to monitor the nodes with the heaviest workloads, which is calculated in operations per second, over the last 24 hours.

Table 102. Activity charts (continued)	
Chart Name	Purpose
Most Active Pools	Use this chart to monitor the pools with the heaviest workloads, which is calculated in operations per second, over the last 24 hours.
Most Active Volumes	Use this chart to monitor the volumes with the highest response times over the last 24 hours.
MDisk Activity	Use this chart to monitor managed disks (M Disks) with the heaviest workloads, which is calculated in operations per second, over the last 24 hours.

File storage systems

The following capacity charts and Inode usage charts are shown for file storage systems:

Table 103. Capacity and Inode charts	
Chart Name	Purpose
Total File System Space	Use this chart to monitor the capacity, the used capacity, and the available capacity of your file systems over the last 30 days.
Capacity by File System	Use this chart to monitor the current capacity and the used and available capacity of your fullest file systems .
Inodes by File System	Use this chart to monitor the availability of Inodes.

Object storage systems

The following capacity charts are shown for object storage systems:

Table 104. Charts for object storage systems	
Chart Name	Purpose
Capacity	Use this chart to monitor the total capacity, used and available capacity of your object storage system.
Sites	Use this chart to monitor the distribution of Accesser and Slicestor nodes across sites.
Capacity by IDA	Use this chart to monitor the used capacity of the vaults that are configured with IDA (Information Dispersal Algorithm).
Failure Tolerance	Monitor the availability and fault tolerance of vaults and Slicestors.

Capacity by pool

Monitor the capacity, fill rate %, and growth rate for your pools. You can also check how much capacity is left before the pools reach their capacity limit.

To help you make decisions about the capacity that you need to store data, monitor the growth in used capacity and the following key capacity metrics:

Adjusted Used Capacity (%)

The amount of capacity that can be used without exceeding the capacity limit.

Example: Adjusted Used Capacity

Before Capacity Limit Was Set 100 GiB Capacity Adjusted Used Capacity = 80% or 80 GiB Available Capacity = 60 GiB Adjusted Used Capacity = 50% or 40 GiB H Capacity-to-Limit = 30% or 40 GiB

The formula for calculating **Adjusted Used Capacity (%)** is (Used Capacity in GiB/Capacity Limit in GiB)*100. For example, if the capacity is 100 GiB, the used capacity is 40 GiB, and the capacity limit is 80% or 80 GiB, then the value for **Adjusted Used Capacity (%)** is (40 GiB/80 GiB)* 100 or 50%. So, in this example, you can use 30% or 40 GiB of the usable capacity of the resource before you reach the capacity limit.

If the used capacity exceeds the capacity limit, the value for **Adjusted Used Capacity (%)** is over 100%.

To add the Adjusted Used Capacity (%) column, right-click any column heading on the Pools page.

See these related values for more information Capacity Limit (%) and Capacity-to-Limit (GiB).

Capacity (GiB)

The amount of capacity that is available for storing data in the pool after formatting and RAID techniques are applied.

Capacity Limit (%) and Capacity Limit (GiB)

The limit that was set on the capacity that is used by your pools. For example, the policy of your company is to keep 20% of the usable capacity of your pools in reserve. So, you log into the GUI as Administrator and set the capacity limit of your pools to 80%.



1. Set the capacity limit



Click the illustration above to find out how to set capacity limits.

The GiB value for the capacity limit for the pool is calculated when you set the value for the **Capacity Limit (%)**.

To add the **Capacity Limit (%)** and the **Capacity Limit (GiB)** columns, right-click any column heading on the **Pools** page.

See these related values for more information **Adjusted Used Capacity (%)** and **Capacity-to-Limit** (GiB).

Zero capacity: When you set the capacity limit for pools, the values shown for **Zero Capacity** are readjusted to take into account the capacity limit of the pool. The date will represent when the capacity limit of the pool is reached. If the pool has already reached the capacity limit, Depleted is

shown. None is shown when a trend in storage consumption can't be detected because the pool's storage isn't being consumed or because not enough data was collected to predict storage consumption.

Capacity-to-Limit (GiB)

The amount of capacity that is available before the capacity limit is reached.



The formula for calculating **Capacity-to-Limit (GiB)** is (Capacity Limit in GiB - Used Capacity in GiB). For example, if the capacity limit is 80% or 80 GiB and the used capacity is 40 GiB, then the value for **Capacity-to-Limit (GiB)** is (80 GiB - 40 GiB) or 40 GiB.

See these related values for more information Capacity Limit (%) and Adjusted Used Capacity (%).

Used Capacity (GiB) and Used Capacity (%)

The amount of usable capacity that is taken up by the data in the pool after data reduction techniques have been applied.

Available for: All storage systems, except FlashSystem A9000 and FlashSystem A9000R.

Recent Fill Rate (%)


The rate at which the capacity of the pool is being consumed over the last 30 days. Use this value to see how quickly your pools are filling up.


The Recent Fill Rate (%) tells you how quickly your pools are filling up


The Recent Fill Rate (%) of the pool is the difference between the fill rate % of the pool 30 days ago and today's fill rate % of the pool.

The fill rate % of the pool is calculated by dividing the capacity of the pool by its used capacity and multiplying it by 100.

Example







The fill rate % of the pool today is $(60/100) \times 100$, which is 60%.

So, in this example, the Recent Fill Rate (%) for the pool is (60% - 40%), which is 20%.

If insufficient historical data is available to calculate the **Recent Fill Rate (%)**, it is the difference between today's fill rate % and the oldest value for the fill rate % that can be calculated in the last 30 days.

Available for: All storage systems, except FlashSystem A9000 and FlashSystem A9000R.

Recent Growth (GiB)

The growth in used capacity over the last 30 days. Use this value to identify the pools with the highest growth in storage consumption.

Recent growth is the difference between the current used capacity for the pool and the used capacity for the pool that was reported 30 days ago. If insufficient historical data is available, it is the difference between the current used capacity for the pool and the oldest value for the used capacity of the pool that was reported in the last 30 days.

Available for: All storage systems, except FlashSystem A9000 and FlashSystem A9000R.

Available Capacity (GiB)

The amount of usable capacity that is not yet used in the pool.

Tier Distribution (%)

For storage systems that support Easy Tier, the distribution of the capacity of the volume extents across each drive class or tier.

Tip: To check the fill rate of all the pools in your storage systems, click **Storage** > **Block Storage Systems**. To check the growth in used capacity, right-click any column heading on the **Block Storage Systems** page and click **Recent Growth (GiB)**.

Capacity by volume

Monitor the provisioned capacity of your volumes to identify the volumes that have sufficient capacity to meet your storage needs.

The following key capacity metrics provide information about storage usage and the availability of capacity for storing data:

Capacity (GiB)

The provisioned capacity of the volume.

Used Capacity (GiB) and Used Capacity (%)

The amount of provisioned capacity that is taken up by written data in a thin-provisioned volume after data reduction techniques are applied.

Unlike thin-provisioned volumes, which use capacity when it is needed, the capacity that is provisioned to standard-provisioned volumes is fully allocated and is no longer available to the pool. That is, the **Used Capacity (%)** of standard-provisioned volumes is reported as Fully Allocated and the **Used Capacity (GiB)** is the same value as the provisioned capacity.

Available for: All storage systems, except FlashSystem A9000 and FlashSystem A9000R.

Available Capacity (GiB)

The difference between the provisioned capacity of the thin-provisioned volume and the used capacity of the thin-provisioned volume.

Because the pool's capacity is dedicated to the standard-provisioned volume, the **Available Capacity** (**GiB**) is reported as 0.00.

Tier Distribution (%)

For storage systems that support Easy Tier, the distribution of capacity for the volume extents across each drive class or tier.

Viewing information about enclosures

View information about enclosures. Enclosures are the metal structures in which various storage components are mounted, including control units, nodes, disks and drives, and power supplies.

Information about enclosures is available for storage systems, including systems that run IBM Spectrum Virtualize.

To view information about enclosures and their related resources, complete any of the following tasks:

Table 105. Viewing information about enclosures			
Actions	Navigation		
View information about the enclosures that are associated with a storage system.	1. In the menu bar, go to Storage > Block Storage Systems.		
Depending on the type of storage system, the following information about an enclosure is shown:	Right-click a storage system and select View Details. Under Internal Resources, click Enclosures.		
Status			
• Type of enclosure: Control, Expansion, or Storage			
Machine Type model (MTM)			
Physical location			
• I/O Group			
 The number of disks, drives, and slots that it contains 			
The number of nodes that it contains			
 The number of power supplies and canisters that it contains 			
The status of its power supplies (PSUs) and canisters			
View the disks in an enclosure.	1. In the menu bar, go to Storage > Block Storage Systems.		
	Right-click a storage system and select View Details .		
	3. Under Internal Resources, click Enclosures.		
	4. In the Disks column, click the number to view more information about each of the disks.		
View the enclosure that a disk or drive is in.	In the menu bar, go to Storage > Block Storage Systems.		
	Right-click a storage system and select View Details .		
	3. Under Internal Resources , click Disks or Drives .		
	4. In the Enclosure column, view the name of the enclosure.		
	5. (Optional) To view information about the enclosure, click its name.		
View the nodes in an enclosure.	In the menu bar, go to Storage > Block Storage Systems.		
	Right-click a storage system and select View Details .		
	3. Under Internal Resources, click Enclosures.		
	4. In the Nodes column, click the number to view more information about each of the nodes.		

Actions	Navigation		
Actions	Navigation		
View the enclosure that a node is in.	1. In the menu bar, go to Storage > Block Storage Systems .		
	Right-click a storage system and select View Details .		
	3. Under Internal Resources, click Nodes.		
	4. In the Enclosure column, view the name of the enclosure.		
	5. (Optional) To view information about the enclosure, click its name.		
View the statuses of power supplies (PSUs) in an enclosure.	In the menu bar, go to Storage > Block Storage Systems.		
	Right-click a storage system and select View Details .		
	3. Under Internal Resources, click Enclosures.		
	4. Right-click an enclosure and select View Properties .		
	5. View the Power Supply Status field.		
View the statuses of canisters in an enclosure.	In the menu bar, go to Storage > Block Storage Systems.		
	Right-click a storage system and select View Details .		
	3. Under Internal Resources, click Enclosures.		
	4. Right-click an enclosure and select View Properties .		
	5. View the Canister Status field.		
Add enclosures to general groups.	In the menu bar, go to Storage > Block Storage Systems.		
Adding storage systems and their associated enclosures to general groups can be helpful when you want to receive alert notifications about	Right-click a storage system and select View Details .		
changes to a group of logically related storage	3. Under Internal Resources, click Enclosures.		
systems.	4. Right-click an enclosure and select Add to General Group .		
	5. Add the enclosure to an existing group or create a new group and click Save .		

Monitoring the status and condition of resources

Monitor the operational condition of storage systems, servers, hypervisors, fabrics, and switches and the status of their internal resources. Use this information to identify potential problem areas in a storage environment.

Table 106. M	Ionitoring the status and condition of resources	
	Explanation	Steps to view status or condition
Status	The status of a resource that is reported by its own hardware to IBM Spectrum Control. For example: • If the cooling fans in a storage system are stopped and the internal temperature is too high, an error status is reported by that storage system. • If a disk on a storage system is starting, a warning status is reported for that disk by the storage system. Tip: The status of internal resources is used to determine the operational condition of the associated top-level resources.	 In the menu bar, go to the resource type that you want to view. For example, if you want to view the status of switches, go to Network > Switches. Right-click a resource and select View Details. A status icon is shown next the image of the resource and its internal resources.
Condition	The overall operational condition of a storage system, server, hypervisor, fabric, or switch. This condition represents the most critical status that was detected on the resource itself and on its internal resources. For example, if an error status was detected on a storage system pool, an error icon is shown for the overall condition of the storage system. If no errors, warnings, or unreachable statuses were detected on a resource or on its internal resources, then a green symbol is shown for the condition of the storage system.	In the menu bar, go to Home > Dashboard. The icons that show overall condition are shown next to the image of a resource type. Resource list pages In the menu bar, go to the resource type that you want to view. For example, if you want to view the condition of switches, go to Network > Switches. The overall condition of a resource is shown in the Condition column, and aggregated in the condition icons at the top of the page.

IBM Spectrum Control provides a number of different icons to help you quickly determine the health of resources.

Table 107. Possible statuses and conditions of resources					
Icon	Icon Health Explanation				
8	Error	A serious problem was detected on a resource or on its internal resources. Resolve these problems as soon as possible.			

Table 107. Possible statuses and conditions of resources (continued)						
Icon	Health	Explanation				
√ 6	Error - Acknowledged	An Error status was detected and acknowledged. An Error - Acknowledged status indicates that a status was reviewed and is either resolved or can be ignored.				
		An acknowledged status is not used when determining the condition of related, higher-level resources. For example, if the status of a volume is Error, the condition of the associated storage system is also Error. If the Error status of the volume is acknowledged, its status is not used to determine the overall condition of the storage system.				
()	Unreachable	A resource is not responding to requests from the IBM Spectrum Control server. This status might be caused by a problem in the network or by a Storage Resource agent that is no longer running and did not communicate that it was shutting down.				
Unreachable - Acknowledged		An Unreachable status was detected and acknowledged. An Unreachable - Acknowledged status indicates that a status was reviewed and is either resolved or can be ignored.				
		An acknowledged status is not used when determining the condition of related, higher-level resources. For example, if the status of a controller is Unreachable, the condition of the associated server is also unreachable. If the Unreachable status of the controller is acknowledged, its status is not used to determine the overall condition of the server.				
<u> </u>	Warning	A Warning status represents potential problems on a resource or on its internal resources. This status is not critical.				
√	Warning - Acknowledged	A Warning status was detected and acknowledged. A Warning - Acknowledged status indicates that a status was reviewed and is either resolved or can be ignored.				
		A Warning - Acknowledged status is not used when determining the condition of related, higher-level resources. For example, if the status of a controller is Warning, the condition of the associated server is also Warning. If the Warning status of the controller is acknowledged, its status is not used to determine the overall condition of the server.				
~	Normal	No warnings or errors were detected on a monitored resource.				
•	Unknown	A resource is known to IBM Spectrum Control but is not monitored. To change an Unknown status, run a probe to collect status information about the resource.				

How the condition of a resource is determined

The operational condition of a resource represents the most critical status that was detected on the resource itself and on its internal resources.

To determine the condition of top-level resources (storage systems, servers, hypervisors, fabrics, and switches), IBM Spectrum Control uses the status of those resources and the status of their internal resources.

The statuses of the following internal resources are used to calculate the overall condition of a top-level resource.

Table 108. Internal resources that are used to determine the condition of top-level resources				
Top-level resource	Internal resources that are used to determine the condition of a top- level resource			
Fabric	Ports Switches			
Switch	■ Blades ✓ Inter-switch connections ■ Logical switches Ports			
Hypervisor	Disks Virtual machines VMDKs			
Server	© Controllers ☑ Disks			

Table 108. Internal resources that are used to determine the condition of top-level resources (continued)					
Top-level resource	Internal resources that are used to determine the condition of a top- level resource				
Block storage system	Disks				
	■ Drives				
	External disks				
	FC ports				
	Host connections				
	I/O groups				
	IP ports				
	Managed disks				
	Modules Modules				
	,Nodes				
	Pools				
	RAID arrays				
	Volumes				
File storage system	Network shared disks				
	■ Nodes				
Object storage system	Network shared disks				
	■ Nodes				
0					

The following statuses of internal resources are used to help calculate the condition of top-level resources:

- Normal
- Warning
- Error

Tip: Statuses that are acknowledged are not used to calculate the overall condition of a top-level resource.

Internal resources for a top-level resource might have different statuses. IBM Spectrum Control uses the most critical status of an internal resource to help determine the overall condition of a top-level resource. For example, in a storage system, a port might have an Error status, a pool might have a Warning status, and multiple controllers might have an Unknown status. In this case, if the storage system itself has a normal status, its overall condition is Error because it is the most critical status that was detected on internal resources.

The following table shows some of the possible combinations of statuses and the resulting, overall condition for a top-level resource.

Error	Unreachable 1	Warning <u> </u>	Normal 	Unknown ²	Resulting condition for a top-level resource
				X	♦ Unknown
			Х		Normal
			Х	Х	Normal
		Х			! Warning
		Х		Х	Warning
		Х	Х	Х	Warning
	Х				S Unreachable
	Х			Х	S Unreachable
	Х		Х	Х	S Unreachable
	Х	Х	Х	X	S Unreachable
Х					8 Error
Х				Х	8 Error
Х			X	Х	& Error
Х		Х	Х	Х	S Error
Х	Х	Х	Х	Х	8 Error

Table 109. Prop	Table 109. Propagation of the statuses for resources (continued)					
Error	Unreachable 1	Warning !	Normal 	Unknown ²	Resulting condition for a top-level resource	

Note:

- 1. The Unreachable status applies only to top-level resources.
- 2. The Unknown status of an internal resources is not used to determine the condition of a top-level resource.

Viewing the overall condition of resources

On the dashboard, you get an overview of the health of your storage resources such as fabrics, hypervisors, servers, storage systems, and switches and of storage consumers such as departments and applications.

The condition of a resource represents the most critical status that was detected on that resource or on its internal resources.

- 1. From the **Home** menu, click **Dashboard**.
- 2. Review the health of your storage environment.

The following conditions icons are shown for block storage systems:

- 🔯 Error
- Unreachable
- A Warning
- Normal

Viewing the condition of specific types of resources

Use resource list pages as a central location for viewing the condition of storage systems, servers, hypervisors, switches, and fabrics.

- 1. In the menu bar, go to the type of top-level resource that you want to view.

 For example, if you want to view the condition of switches, go to **Network** > **Switches**.
- 2. To view the total number of different conditions that were detected for the type of resource, check the condition icons on the resource list page.

The number next to the icon represents how many occurrences of each condition were detected. For example, if 15 switches are monitored, but five switches have internal resources with a warning status, two switches have internal resources with an unreachable status, and three switches have internal resources with an error status, the following condition information is displayed:



5 Normal



5 Warning



2 Unreachable



- 3. To view the condition for specific resources, check the **Condition** column in the list of resources.
- 4. Click the **Search** icon \(\bigcirc\) to filter the list so that only resources with a specified condition are shown. You can use filtering to identify the resources with the most serious conditions in your environment.
- 5. Select Condition.
- 6. In the **Choose Value** list, select the condition on which you want to filter.

The list of resources is automatically refreshed and only resources with the filtered status are shown.

For example, select Warning



to view only the resources with a Warning condition.

Viewing the status of resources

Use detail pages to view the status of storage systems, servers, hypervisors, switches, and fabrics and the status of their internal resources. The status of a resource is reported to IBM Spectrum Control by its hardware.

The status of a resource is different than the operational condition that is calculated by IBM Spectrum Control for top-level resources:

- Condition is determined by propagating the statuses of resources and is only shown for top-level resources.
- *Status* represents the status of a resource as reported by its hardware and is shown for top-level resources and internal resources. The following examples illustrate statuses that might be reported by a resource:
 - If the cooling fans in a storage system are stopped and the internal temperature is too high, an error status is reported by that storage system.
 - If a disk on a storage system is starting, a warning status is reported for that disk by the storage system.

Tip: View condition to identify which top-level resources might be encountering problems. View status to identify the specific resources that are causing the problems in a top-level resource.

- In the menu bar, go to the resource type that you want to view.
 For example, if you want to view the status of block storage systems and their internal resources, go to Storage > Block Storage Systems.
- 2. Right-click a resource and select View Details.
- 3. To view the status of the top-level resource, view the status icon that is displayed next to its image on the page.
- 4. To view the aggregated status of an internal resource, view the status icon that is displayed next to its image in the **Internal Resources** section.

If a resource has a status other than normal, an icon is displayed for the most critical status.

For example, if a server has three controllers with an error status and two controllers with a warning status, the error status icon is shown for controllers in the **Internal Resources** section:



5. To view the status of a specific internal resource, click the name of that internal resource in the **Internal Resources** section.

A list of the monitored internal resources is displayed. For example, on the details page for a server, click **Controllers (5)** to view a list of the five controllers that are associated with the server.

6. Check the **Status** column to view the status for each internal resource.

Tip: When you view information about some internal resources, the Status column might show values that are more specific than Error, Unreachable, Warning, Unreachable, Unknown, and Normal. Use this additional status to determine the cause of the problem for an internal resource.

For example, in the **Internal Resources** section on the details page for a storage system, an error status might be shown for managed disks:



When you view the list of managed disks, the Status column might show an Offline status for a specific managed disk. When a managed disk is reported as offline, an error status icon 🍑 indicates that status.

Acknowledging the condition and status of resources

Sometimes the status of resources might represent problems that commonly occur but can be ignored. In such cases, you can acknowledge those statuses so that they are not used when determining the overall condition of storage systems, servers, hypervisors, fabrics, and switches.

Acknowledging the condition of top-level resources

A summary of the conditions for storage systems, servers, hypervisors, fabrics, and switches is available in the Dashboard view and on resource list pages. You can acknowledge the condition of a top-level resource so that its condition is not shown as Error, Unreachable, or Warning on those pages.

- 1. In the menu bar, go to the type of resource that has a condition you want to acknowledge. For example, if you want to acknowledge the condition of a switch, go to **Network > Switches**.
- 2. In the list of resources, locate the resource with the Error, Unreachable, or Warning condition that you want to acknowledge.
- 3. Right-click the resource and select **Acknowledge Condition**. A window shows that the condition for the specified resource was acknowledged.
- 4. Optional: To acknowledge the condition of multiple resources at the same time, press Ctrl and click each resource. To select a series of resources, select the first resource, and then press Shift and click the last resource. When the resources are selected, right-click a row in the list and select **Acknowledge Condition.**
- 5. Click Close.

When you acknowledge the condition of a top-level resource, the following actions are taken:

Dashboard view

- The condition of the acknowledged resource is not used to determine the condition icon that is displayed for that resource type on the Dashboard view. For example, if you acknowledge the unreachable condition for a switch, the condition of that switch is not used to determine the number of unreachable conditions that is shown for switches on the Dashboard view.
- The acknowledged condition is shown as **Acknowledged** for a resource on the Dashboard view. For example, if 10 switches have Unreachable conditions and you acknowledge the condition for one of them, the following information is shown for switches:



9 Unreachable (1 Acknowledged)

Resource details page

• The condition icons at the top of the page are updated. For example, if 10 switches have Unreachable conditions, and you acknowledge one of them, the total number Unreachable conditions decreases by one and the number of Unreachable - Acknowledged conditions increases by one. The following information is updated at the top of the **Switches** page:





• The value in the **Condition** column for the resource is updated to show an acknowledged condition. For example, if you acknowledge the Unreachable condition for a switch, the following value is shown in the **Condition** column:



Related reference

"Resources that you can monitor" on page 111

You can use IBM Spectrum Control to monitor storage systems, servers, hypervisors, fabrics, and switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

"How the condition of a resource is determined" on page 248

The operational condition of a resource represents the most critical status that was detected on the resource itself and on its internal resources.

Acknowledging the status of internal resources

The statuses of internal resources are used to determine the condition of the associated storage systems, servers, hypervisors, fabrics, and switches. Sometimes these statuses might represent problems that commonly occur or are known, but can be ignored. In such cases, you can acknowledge the status of an internal resource so that it is not used to determine the condition of a top-level resource.

For example, if the status of a volume is Error, the condition of the associated storage system is also Error. If the Error status of the volume is acknowledged, its status is not used to determine the overall condition of its storage system. In this case, if the other internal resources of the storage system are Normal, then the condition of the storage system is also Normal.

- 1. In the menu bar, go to the type of resource in which the internal resource is located. For example, if you want to acknowledge the status of internal resources for a switch, go to **Network** > Switches.
- 2. Right-click a resource and select View Details. For example, if you want to acknowledge the status of a port on a switch, right-click that switch and select View Details.
- 3. In the Internal Resources section of the resource details page, click the name of the internal resource.
 - For example, if you want to acknowledge the status of a port on a switch, click Ports.
- 4. Right-click the resource and select **Acknowledge Status**.
 - A window shows that the status for the specified resource was acknowledged. For example, if you want to acknowledge the status of a port, right-click the port and select Acknowledge Status.
- 5. Optional: To acknowledge multiple resources at the same time, press Ctrl and click each resource. To select a series of resources, select the first resource, and then press Shift and click the last resource. When the resources are selected, right-click a row in the list and select **Acknowledge Status**.
- 6. Click Close.

The status of the related, top-level resource might take several minutes to update.

When you acknowledge the status of an internal resource, the following actions occur:

- The status of the internal resource is no longer used to determine the condition of the associated toplevel resource.
- The status icons at the top of the page for the internal resource are updated.

For example, if you acknowledge the Error status for a controller on a server, the total number of Error statuses decreases by one and the number of Error - Acknowledged statuses increases by one on the **Controllers** page. If originally there were five controllers with Error statuses, the following updated status information is shown:



🔯 4 Error 蜷 1 Error - Acknowledged

Related reference

"Resources that you can monitor" on page 111

You can use IBM Spectrum Control to monitor storage systems, servers, hypervisors, fabrics, and switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

"How the condition of a resource is determined" on page 248

The operational condition of a resource represents the most critical status that was detected on the resource itself and on its internal resources.

Monitoring vaults in IBM Cloud Object Storage

Monitor access risk and storage risk for vaults in IBM Cloud Object Storage.

You can use this information to complete the following tasks:

- Identify the vaults that cannot be accessed and the vaults that are at risk of access failure.
- Identify the vaults that are not available for write operations and the vaults that are at risk of storage failure.

Monitoring access to vaults in IBM Cloud Object Storage

In these scenarios, you use IBM Spectrum Control to identify the vaults that cannot be accessed and the vaults that are at risk of access failure. Then, you identify the access pools and the COS Accesser nodes that the vaults depend on.

How the accessibility value for a vault is determined

The accessibility value that is calculated for each vault comprises the number of COS Accesser nodes that are available to access the vault. For example, if five COS Accesser nodes are configured for a vault and four of the nodes are in a failed state, then the accessibility value is one.

On the details page for IBM Cloud Object Storage, the Vaults by Accessers chart shows the vaults that are inaccessible and the accessibility values for the remaining vaults in the storage system.

Investigating vaults that cannot be accessed

In this scenario, you use IBM Spectrum Control to identify the vaults that cannot be accessed because of COS Accesser node failures.

You have a backup application that uses IBM Cloud Object Storage vaults. Because the application is critical to your business, you want to quickly identify vaults that cannot be accessed because of COS Accesser node failures. You also want to know the access pools that the vaults depend on and the COS Accesser nodes that failed in those pools.

- 1. To determine whether any of your vaults cannot be accessed, click **Storage > Object Storage Systems**.
- 2. Right-click the storage system that is used by your backup application and click **View Details**.
- 3. In the navigation pane, click **Overview**.

 The first line of the Vaults by Accessers chart shows the number of vaults that cannot be accessed, for example, Inaccessible 4.
- 4. To identify the vaults that are inaccessible, click **Vaults** in the navigation pane. To sort the data, click the **Accessibility** column.

A value of 0 in the Accessibility column indicates that the vault is inaccessible. The error status indicates that the vault is inaccessible because none of the COS Accesser nodes that are configured for the vault are available. In this scenario, all the inaccessible vaults are configured to use just one access pool. The name of that access pool is shown in **Access Pools**.





5. To identify the COS Accesser nodes that failed in the access pool, click the access pool name. Then, click the **Accesser Nodes** tab.

To restore access to the vaults, fix the COS Accesser nodes that failed or add COS Accesser nodes to the access pool.

Monitoring vaults that are at risk of access failure


In this scenario, you use IBM Spectrum Control to check which vaults are at risk because of COS Accesser node failures.

You have a backup application that uses IBM Cloud Object Storage vaults. Because the application is critical to your business, you want to monitor the tolerance of the vaults to COS Accesser node failures and identify the vaults that are at risk of access failure.

For the vaults with a high risk of failure, you want to identify the access pools that the vaults depend on and the status of the COS Accesser nodes in the access pools.

- 1. Click Storage > Object Storage Systems.
- 2. Right-click the storage system that is used by your backup application and click View Details.
- 3. In the navigation pane, click **Overview**.
 - The Vaults by Accessers chart shows the tolerance of the vaults to COS Accesser node failures. For example, a line on the chart that shows a value of 1 Node Failure 57 Vaults at Risk indicates that 57 vaults rely on a single COS Accesser node for access. If the node fails, the vault cannot be accessed.
- 4. To identify the vaults with an accessibility value of 1, click **Vaults** in the navigation pane. To sort the data, click the **Accessibility** column.






In this scenario, there are some vaults that depend on only one access pool. The name of that access pool is shown in **Access Pools**. For other vaults that depend on multiple pools, a number is shown in **Access Pools** instead.


Tip: To reduce the risk of access failure for the vaults that depend on only one access pool, add COS Accesser nodes to the access pool or deploy the vaults to other access pools.

For the vaults that depend on multiple access pools, the warning status the COS Accesser nodes in the pools are not available.



5. To investigate the cause of the access warning, click the number in the Access Pools column.

You see that one of the access pools is shown with an error status , which indicates that the COS Accesser nodes in the pool are not available.



6. To identify the COS Accesser nodes that failed in the access pool, click **Accesser Nodes** in the navigation pane. To sort the data, click the **Access Pool** column.

Fix the COS Accesser nodes that failed or add COS Accesser nodes to the access pool.

Calculating the failure tolerance for vaults

The failure tolerance value for vaults represents the tolerance of a vault to drive failures across all of the COS Slicestor nodes in its storage pool. The lower the value for drive failure tolerance, the greater the risk that write operations can no longer be completed for the vault.

The calculation of drive failure tolerance is based on the following values:

- The write threshold value in the IDA that is configured for the vault. The write threshold is the minimum number of COS Slicestor nodes that must be available to complete write operations for the vault.
- The number of available COS Slicestor nodes in the storage pool that the vault belongs to.
- The value for drive failure tolerance for each of the available COS Slicestor nodes in the storage pool. This value represents the tolerance of the node to drive failures and is based on the number of failed drives in the node and the drive error threshold. A value of 2 means that the node will be unavailable to store vault data if 2 drives fail in the node.

For example, a vault in your IBM Cloud Object Storage environment is configured with an IDA of 8-4-6. The IDA determines that the vault data is stored across 8 COS Slicestor nodes and a minimum of 6 nodes must be available to complete write operations for the vault.

The nodes have the following values for drive failure tolerance:

Node	Drive Failure Tolerance
node_1	3
node_2	4
node_3	0*
node_4	2
node_5	4
node_6	1
node_7	3
node_8	3

^{*}A failure tolerance value of 0 means that node_3 is not available to store vault data.

7 COS Slicestor nodes are available in the storage pool and a minimum of 6 nodes must be available to complete write operations. Therefore, if 2 more nodes fail, the write threshold value is exceeded and write operations can no longer be completed for the vault.

Use the following steps to calculate the value for drive failure tolerance for a vault:

- 1. Use the IDA column on the **Vaults** page to identify the write threshold value for the vault. In the example, the write threshold value is 6.
- 2. Use the **Slicestor Nodes** page to determine the number of available COS Slicestor nodes in the storage pool that the vault belongs to. A value > 0 in the Drive Failure Tolerance column means that the node is available. In the example, 7 nodes are available.
- 3. Calculate how many nodes must fail for the number of available nodes to fall below the write threshold value. In the example, (7 6) + 1 = 2. When 2 nodes fail, the vault is not available for write operations.
- 4. Use the Drive Failure Tolerance column on the **Slicestor Nodes** page to identify the 2 available nodes in the storage pool that have the lowest tolerance to drive failures.

In the example, the 2 nodes are as follows:

Node	Drive Failure Tolerance
node_4	2
node_6	1

5. To calculate the overall failure tolerance for the vault, sum the failure tolerance values for the 2 nodes. 2 + 1 = 3.

The drive failure tolerance for the vault is 3. If 3 drives fail, write operations can no longer be completed for the vault.

Monitoring the performance of resources

IBM Spectrum Control can collect information about the performance of storage systems and switches. This information includes key performance metrics that can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage.

You can use the performance monitoring to complete the following tasks:

- Measure, compare, and troubleshoot the performance of switches, storage systems, and their internal resources.
- Review the alerts that were triggered when the performance if a resource fell outside of a specific range.
- Customize views of performance so that you can analyze specific resources and metrics during time ranges that you specify.
- View performance information in a chart or table format to help you quickly identify where and when performance issues are occurring. The chart is a visual representation of how the performance of resources trend over time.
- Drill down into resources to view detailed information about the performance of internal and related resources. For example, if a SAN Volume Controller is shown in the chart, you can quickly view and compare the performance of its internal and related resources, such as disks, volumes, ports, managed disks, and back-end storage.
- Implement server-centric monitoring of SAN resources without requiring a Storage Resource agent. When you add an agentless server, IBM Spectrum Control correlates the server with the known host connections on monitored storage systems. If matches are found, you can view the amount of storage that is assigned to the server, and trace that storage back to the storage system. You can then view details about the internal resources of the related storage system, including performance information.
- Export performance information to a CSV file. A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application.

Before you begin

Before you can view performance information for resources, you must complete the following tasks:

Task	Learn more
Add storage systems and switches for monitoring by IBM Spectrum Control.	For information about how to add these resources, see the following topics: • "Adding storage systems" on page 118 • "Adding fabrics and switches" on page 129
Ensure that performance data is being collected for a resource. Typically, data collection is scheduled when resources are added for monitoring.	For information about verifying data collection and scheduling performance monitors, see the following topics: • "Verifying that a performance monitor is running for a resource" on page 145 • "Creating performance monitors in IBM Spectrum Control" on page 144

Task	Learn more
Define performance alerts to be notified if the performance of a resource falls outside a specified range and might represent a potential problem.	For information about how to define performance alerts, see "Defining alert definitions for performance changes" on page 162.
When you define a performance alert for a resource, select a specific metric that you want to measure. For example, you can define a threshold that notifies you when the total I/O port rate for a storage system falls outside a specified range.	

Related tasks

Tutorial: Identifying the source of slow drain problems caused by depletion of buffer credits
Use this tutorial to find out how to use IBM Spectrum Control to identify a host that has depleted buffer credits that are causing a slow drain condition.

Related information

[Video] IBM Spectrum Control Performance Troubleshooting Techniques

Viewing performance information

View the performance of storage systems, switches, and their internal resources. You can also view the performance of resources that violated a performance threshold and generated an alert.

Before you view the performance of resources, ensure that performance data was collected for those resources during the time ranges that you want to analyze. IBM Spectrum Control uses performance monitors to collect metrics for measuring the performance of storage systems and switches. For information about how to verify that performance monitors were run, see "Verifying that a performance monitor is running for a resource" on page 145.

You can access the performance view for a resource from different locations in IBM Spectrum Control:

Task	Accessing the performance view in the GUI
View the performance of specific storage systems, switches, and their internal resources.	1. In the menu bar, go to the type of resource that you want to view. For storage systems, depending on the type of storage system that you want to view, go to Storage > Block Storage Systems, Storage > File Storage Systems, or Storage > Object Storage Systems. For switches, go to Network > Switches.
	Right-click a resource and select View Performance . By default, the following information is displayed in a chart for each type of resource:
	 For storage systems, the top 5 storage systems with the highest total I/O rate are shown.
	For switches, the top 5 switches with the highest total port data rate are shown.
	3. To view the performance of an internal resource, right-click the top-level resource in the chart legend and select the option for the internal resource that you want to view. For example, for a SAN Volume Controller, right-click its name in the chart legend and select MDisk Performance to view the performance of its managed disks.
	For managed disks, the top 5 contributors to the storage system workload are shown on the chart. The default metrics are total back-end I/O rate and total back-end response time.
	Tips:
	To add or change the metrics that are displayed
	for a resource, click the select metrics icon and select the metrics to display.
	You can also view the performance of internal resources on the details page for the related, top-level resource. For more information, see Table 110 on page 264.
	To view the performance of multiple resources at the same time, press Shift or Ctrl and click the resources that you want to view. Press Shift and click to select consecutive rows in the list of resources; press Ctrl and click to select nonconsecutive rows. Then, right-click any of the selected resources and select View Performance. Each resource is represented by a separate line in the performance chart and separate rows in the table view.

Task	Accessing the performance view in the GUI	
View the overall performance of resources of the same type.	1. In the menu bar, go to the type of resource that you want to view. For storage systems, depending on the type of storage system that you want to view, go to Storage > Block Storage Systems, Storage > File Storage Systems, or Storage > Object Storage Systems. For switches, go to Network > Switches.	
	Click the Performance tab. By default, the following information is displayed in a chart:	
	For storage systems, the top 5 storage systems with the highest total I/O rate are shown.	
	For switches, the top 5 switches with the highest total port data rate are shown.	
	Tip: To add or change the metrics that are displayed for a resource, click the select metrics	
	icon 🖶 and select the metrics to display.	
View the alerts that were generated by threshold violations and the performance of the resources where the violation occurred.	For information about how to view performance alerts, see "Viewing performance alerts" on page 275.	
View the performance of a resource that is included in a specific performance monitor.	In the menu bar in the web-based GUI, go to Home > Performance Monitors.	
	In the Name column, locate the name of the resource that you want to view.	
	3. Right-click the performance monitor and select View Performance .	
View predefined reports and create custom reports that show the performance of resources. In the optional Cognos Analytics reporting tool, you can	1. Go to the URL for your Cognos Analytics server. The URL is similar to this URL: http://myhostname:9300/bi	
view reports about storage systems, components of storage systems, switches, and switch ports.	2. Click Team Content in the Welcome portal.	
or storage systems, switches, and switch ports.	To view a predefined performance report, complete the following steps:	
	a. Click IBM Spectrum Control Predefined Reports .	
	b. Navigate to the performance report that you want to view and click the name of the report.	
	To create custom performance reports, complete the following steps:	
	a. Click IBM Spectrum Control Packages .	
	b. Right-click Performance , then click Create report.	
	c. Click a template for the report.	

You can also view the performance of internal resources on the details pages for the related, top-level resources. For example, you can view the performance of the controllers for a storage system on the details page for that storage system. To access the performance view for an internal resource on a details page, select the type of internal resource and click the **Performance** tab.

The following information summarizes the internal resources and locations where you can access a performance view.

Table 110. Internal resources and locations where you access the performance view		
Internal resource	Locations in the GUI where you can access the performance view	
Pools	 Storage > Pools Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click Pools 	
	Available only for pools that are associated with the following storage systems and virtualizers: System Storage DS series, ESS, SAN Volume Controller, the IBM Storwize family, IBM Spectrum Accelerate, and XIV systems.	
Volumes	 Storage > Volumes Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click Volumes 	
Controllers	Click Storage , and then click Block Storage Systems , File Storage Systems , or Object Storage Systems , depending on the storage system that you require. Right-click the storage system, click View Details , and click Controllers	
Disks	Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click Disks Available only for disks that are associated with storage virtualizers such as SAN Volume Controller and the IBM Storwize family.	
Host Connections	Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click Host Connections	
	Available only for host connections that are associated with the following storage systems and virtualizers: System Storage DS series, ESS, SAN Volume Controller, the IBM Storwize family, and XIV systems.	
I/O Group	Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click I/O Groups	

Table 110. Internal resources	and locations where you access the performance view (continued)	
Internal resource	Locations in the GUI where you can access the performance view	
Modules	Click Storage , and then click Block Storage Systems , File Storage Systems , or Object Storage Systems , depending on the storage system that you require. Right-click the storage system, click View Details , and click Modules	
Nodes	Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click Nodes	
Managed Disks	Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click Managed Disks	
	Available only for nodes that are associated with storage virtualizers such as SAN Volume Controller, the IBM Storwize family, Hitachi VSP, and NetApp storage systems running ONTAP 9.	
Ports on storage systems	Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click Ports	
Ports on switches	Network > Switches, and click Ports	
RAID Arrays	Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click RAID Arrays	
	Available only for RAID arrays that are associated with System Storage DS series and ESS storage systems.	

Related information

Performance statistics (for SAN Volume Controller systems)
Performance statistics (for Storwize V7000 Unified systems)

How performance information is displayed

The performance view is displayed when you view the performance of resources, which includes resources in alerts and performance monitors. Information in the performance view is organized into two main sections: a chart or table and a legend.

Performance chart

The top section of the performance view shows information about the selected resources. You can view this information in the following formats:



The chart shows a visual representation of how the performance of a resource trends over time. Each line on the chart represents a metric and a resource. For example, if you select two metrics and three resources, six lines are shown on the chart. The y-axis shows the unit of measurement for a metric. If more metrics were selected with a different unit of measurement, an extra y-axis is shown on the right side of the chart window.

Table

The table shows performance and asset information that is formatted into rows and columns. Each row represents a resource and a time stamp from the chart; each column represents a metric from the chart. For example, if two metrics and three resources are displayed on a chart, and each of the lines on the chart have 10 data points, 30 rows and 2 columns are shown in the table.

To view other metrics and asset information for a resource, right-click anywhere in the header row of the table and select extra columns. The type of resource determines the metrics and information that is available.

Performance chart legend

The bottom section of the performance view shows more information about the resources in the chart. The information is formatted into rows and columns. Each row represents a resource that was selected for the view. Each column provides asset and performance information about a resource.

Each metric on the chart is represented by two columns in the legend. The first column includes the weighted average of the metric for each resource in the selected time range. The second column represents the peak data point of the metric for each resource in the selected time range. For example, if a metric in the chart is Total I/O Rate, the following columns are displayed in the legend: Avg Total I/O Rate, Worst Total I/O Rate. The worst value can be the minimum or the maximum number, depending on the metric. For example, the worst value for an I/O rate metric is a maximum number, while the worst value for the Cache Hit Percentage metric is a minimum number.

You can also view information about the resources that are related to the resources in the legend. For example, if a SAN Volume Controller is listed in the legend, right-click it to view the performance of its internal and related resources, such as disks, volumes, ports, managed disks, and back-end storage. For more information about the resources that you can view in the chart legend, see "Resources in the chart legend" on page 269.

Tip:

- If you hide a resource in the performance chart, the row for that resource remains visible in the chart legend.
- In the chart legend, information about a resource represents the state of that resource during the selected time range. Specifically, this information does not represent the current state of a resource, but instead shows the final state of the resource during the selected time range.

For example, if the current date is January 1, and you view a volume with the time range set to December 1 to December 7, the legend shows the capacity of that volume on December 7. If the capacity of the volume was changed between December 7 and January 1, this historical capacity is different from the current capacity of the volume on January 1.

Controls

Each performance view includes controls for customizing how information is displayed. The type of view and the resources that you are viewing determine which controls are available.



View chart

View performance information in a chart format. The chart shows a visual representation of how resource performance trends over time. Each line on the chart represents a metric and a resource. The y-axis shows the unit of measurement for a metric. If more metrics were selected with a different unit of measurement, an extra y-axis is shown on the right side of the chart window.

Hover the mouse pointer over points on a line to view a snapshot of performance information at a specific time.



View table

View performance information in a table format. Each row represents a resource and a time stamp. Each column represents a metric. You can view other metrics and information for a resource by right-clicking anywhere in the header row for the table and selecting more columns. The type of resource determines the metrics and information that is available.

For information how to filter, sort, and customize the columns in a table, see "Customizing lists" on page 454.



Metrics

Add metrics to the performance chart. On the **Select Chart Metrics** dialog, metrics are organized into Volume, Disk, Pool, Port, and Node categories. The metrics and categories that are available depend on the type of resource that is being shown in the chart. The number next to the name of the category represents the number of metrics that are currently selected from that category.

You can select multiple metrics at the same time and from different categories, but you cannot include more than two unit types in the same view. For example, if you select metrics that use % and ops/s as units of measurement, you cannot select more metrics that use different units of measurement such as KiB/op or MiB/s.

For a list of metrics that you can view for resources, see "Performance metrics" on page 277.



Hide and show resources

When the chart includes multiple resources, you can click the icon next to a resource to show only the line for that resource. Each icon is shown in a different color to match the color of line for the resource.

You can also show and hide resources in the chart by selecting resources in the chart legend. To select multiple resources at the same time, press Shift or Ctrl and click those resources. Press Shift and click to select consecutive rows in the legend; press Ctrl and click to select non-consecutive rows.

You can show up to 10 resources in a chart at the same time.





Specify granularity

Determine the granularity of the data that is shown in a performance view. Granularity determines the points that are shown on the y-axis of a chart, the points on the lines in a chart, and the rows in a table. When you first view performance information, the default time range is the last 12 hours and the granularity matches the frequency of data collection by the performance monitor (sample).

To change the granularity of chart or table, select one of the following options from the granularity menu:



(daily)

Set the granularity to show 1-day increments for performance information that is shown on the chart and table.

This granularity is only available if the time range of the chart is more than 1 day.



(hourly)

Set the granularity to show one-hour increments for performance information that is shown on the chart and table.

This granularity is only available if the time range of the chart is more than 1 hour.



(Sample)

Set the granularity to match the frequency of data collection by a performance monitor. For example, if a performance monitor is scheduled to collect sample data every 15 minutes, each point on the chart is shown in 15-minute increments.

This granularity is only available if the time range of the chart is less than 2 weeks.

If the performance monitor collects sample data at 1-minute intervals, data is displayed as 1-minute intervals in the chart only when data is available for all of the selected time range. If data is not available at 1-minute intervals for any portion of the selected time range, the data is displayed in 5-minute intervals.

For example, the performance monitor for a resource collects sample data every 5 minutes. Six hours ago, the performance monitor was changed to start collecting data every minute. You want to view sample data for the resource for the last 12 hours. However, the performance monitor did not collect sample data every minute for 6 of the last 12 hours. Therefore, the data is displayed in 5-minute intervals rather than 1-minute intervals.

Restriction: Each granularity has a maximum time range for displaying data. This time range is dynamically calculated based on the resources and metrics that are shown on the chart, and the data points per line. When the maximum range is exceeded, the granularity icon is disabled and a message indicates that there are too many lines on the chart. To avoid this problem, reduce the time range or the number of resources or metrics that are selected.

Specify a time range

The time range of the performance information is shown below the chart. You can change this range to display information for different times when data was collected. When you first access the performance view, the default time range is the last 12 hours.

Ensure that performance data was collected during the time range that you select. If data was not collected during the time range, the chart and table are blank. If data collection was interrupted during the time range, the chart and table show gaps for the time increments when data was not collected.

For example, if you select a time range for the last 7 days, but data was not collected on days 4 and 5, the lines in the chart do not show data for days 4 and 5.



Hide and show chart controls

The performance view lists the names of resources that are being shown in the chart and the controls for customizing that chart. Click the left arrow icon to hide the controls and display the chart in the entire window. Click the right arrow icon again to show the controls.



Export information about the chart to a file

Export information on a performance view to a CSV file.



Open the performance view in a separate web browser window

Open a duplicate of the current performance view in a separate web browser window. You can change the information that is displayed in this separate window while retaining the original performance view for comparison.



Synchronize the time range across all the open performance views

Synchronize the time range across all the performance views that are displayed in separate browser windows. Use this action when you change the time range in a performance view and want to apply the same time range to the other performance views.

Restriction: This action affects only the performance views that are displayed in separate browser windows. The time range of the performance view that is shown as part of the main window for the GUI is not affected.

For example, you can view the storage system volumes that are assigned to a server. You can then open separate browser windows for the host connections, pools, and managed disks that are related to one of the volumes. If you change the time range in the performance view for host connections, click the synchronize button to apply the same time range to the views for pools and managed disks. The performance view of the original volume is not changed.

Actions for managing the resources in the chart legend

The chart legend in bottom section of the view shows more information about the selected resources. This information is organized into rows and columns, where each row represents a resource.

When you select one or more resources in the legend, the following actions are available in the **Actions** menu:

View Properties

View key details about a resource, including asset, status, configuration, capacity, and performance information.

View resource Performance

View the performance of resources that are internal or related to a resource in the chart legend. For example, when you view the performance of a SAN Volume Controller, you can right-click it and view the performance of its internal resources, such as disks, volumes, ports, managed disks, and back-end storage.

Information about an internal or related resource is shown in a separate web browser window. This window uses the same time range as the performance view in the main window of the GUI.

Restriction: If performance data was not collected for an internal or related resource, that resource is not shown in the view. For example, if you right-click a storage system and select **Volume performance**, only the volumes for which performance data was collected are shown.

For more information about the resources that you can view in the chart legend, see <u>"Resources in</u> the chart legend" on page 269.

Key Metrics View

View a set of predefined metrics for resources in the chart legend.

Custom View

Customize the set of metrics that are shown for resources in the chart legend.

Resources in the chart legend

In the performance view, performance metrics and related information for resources are shown in a chart and in the chart legend. For resources in the chart legend, you can open separate performance views for their internal or related resources. The separate performance views use the same time range as the performance view in the main window of the GUI.

For example, when you view the performance of a SAN Volume Controller, you can right-click it in the chart legend to open a performance view for its internal and related resources, such as disks, volumes, ports, managed disks, and back-end storage.

The type of resource determines the internal and related resources that you can view. The following tables show the internal and related resources that are available when you right-click a resource in the chart legend.

Table 111. Resources that you can view for DS6000, DS8000, and ESS storage systems in the chart legend		
Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Storage system	Volume, RAID array, pool, controller, host connection, port	Storage virtualizer, switch
Volume	Storage system, RAID array, pool, controller, host connection	Managed disk
RAID array	Storage system, volume, pool, controller, host connection	None
Pool	Storage system, volume, RAID array, controller, host connection	None
Controller	Storage system, volume, RAID array, pool, host connection	None
Host connection	Storage system, volume, RAID array, pool, controller	None

Table 111. Resources that you can view for DS6000, DS8000, and ESS storage systems in the chart legend (continued)			
Right-click the resource in the chart legend to view the performance of the following resources Right-click the resource in the chart legend to view the performance of the following related resources			
Storage system port Storage system Switch port			

Table 112. Resources that you can view for SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified storage systems in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Storage system	Volume, pool, node, I/O group, disk, managed disk, host connection, port	Back-end storage system, storage virtualizer, switch
Volume	Storage system, pool, node, I/O group, disk, managed disk, host connection	None
Pool	Storage system, volume, node, I/O group, disk, managed disk, host connection	None
Node	Storage system, volume, pool, I/O group, disk, managed disk, host connection, port	None
I/O group	Storage system, volume, pool, node, disk, managed disk, host connection, port	None
Disk	Storage system, volume, pool, I/O group, node, managed disk, host connection	None
Managed disk	Storage system, volume, pool, node, I/O group, disk, host connection	Back-end volume
Host connection	Storage system, volume, pool, node, I/O group, disk, host connection	None
Storage system port	Storage system, I/O group, node	Switch port

Table 113. Resources that you can view for IBM Spectrum Scale in the chart legend			
Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources	
Storage system	File system, node	None	
File system	Storage system, node	None	

Table 113. Resources that you can view for IBM Spectrum Scale in the chart legend (continued)		
Right-click the resource in the chart legend to view the performance of the following resources Right-click the resource in the chart legend to view the performance of the following related resources		
Node	Storage system, file system	None

Table 114. Resources that you can view for IBM Spectrum Accelerate in the chart legend		
Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources Right-click the resource in the chart legend to view the performance of the following related resources	
Storage system	Module, pool, volume, host connection, port	Storage virtualizer, switch
Module	Storage system, pool, volume, host connection, port	None
Pool Storage system, module, volume, None host connection		None
Volume	Storage system, module, pool, host connection	Managed Disk

Table 115. Resources that you can view for XIV systems in the chart legend		
chart legend to view the performance of the following		Right-click the resource in the chart legend to view the performance of the following related resources
Storage system	Module, pool, volume, host connection, port	Storage virtualizer, switch
Module	Storage system, pool, volume, host connection, port	None
Pool	Storage system, module, volume, host connection	None
Volume	Storage system, module, pool, host connection	Managed Disk
Host connection	Storage system, module, pool, volume	None
Storage system port	Storage system, module	Switch port

Table 116. Resources that you can view for IBM Spectrum Accelerate in the chart legend		
Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources Right-click the resource in the chart legend to view the performance of the following related resources	
Storage system	Module, pool, volume, host connection, port	Storage virtualizer, switch

Table 116. Resources that you can view for IBM Spectrum Accelerate in the chart legend (continued)		
chart legend to view the chart legend to view the		performance of the following
Module	Storage system, pool, volume, host connection, port	None
Pool	Storage system, module, volume, host connection	None
Volume	Storage system, module, pool, host connection	Managed Disk

Table 117. Resources that you can view for SMI-S BSP systems in the chart legend		
chart legend to view the chart legend to view the		performance of the following
Storage system	Volume, port	Storage virtualizer, switch
Volume Storage system Managed disk		
Storage system port	Storage system	Switch port

Table 118. Resources that you can view for switches in the chart legend		
chart legend to view the chart legend to view the		performance of the following
Switch	Port	Storage system, switch
Port	Switch	Storage system port, switch port

Table 119. Resources that you can view for servers in the chart legend		
Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Server	None	Storage system, volume, pool, controller, module, node, I/O group, disk, managed disk, host connection, switch, storage system port, switch port

Saving URLs for performance views

You can create a bookmark or a favorite for a performance view.

The link in the bookmark or favorite stores details of the performance view, such as the resources, metrics, interval, and other details. The link also stores the time period that is specified in the time selector at the top of the chart.

You can also copy the URL from the address bar if you want to share the URL with colleagues.

Do one of the following steps:

- In Firefox, right-click in the performance view and click **Bookmark This Page**.
- In Internet Explorer, right-click in the performance view and click **Add to favorites**.

Identifying performance issues for IBM Spectrum Virtualize storage systems

You can identify key metrics that are outside of a standard range for resources that run IBM Spectrum Virtualize by using the key performance indicators charts.

Before you view the performance of your resources that run IBM Spectrum Virtualize, ensure that performance data was collected for those systems during the time ranges that you want to analyze. IBM Spectrum Control uses performance monitors to collect data so that you can measure the performance of storage systems and switches. To verify that performance monitors were run, follow the instructions in Verifying that a performance monitor is running for a resource.

In this documentation, IBM Spectrum Virtualize is used to refer collectively to IBM SAN Volume Controller, IBM Spectrum Virtualize for Public Cloud, IBM Spectrum Virtualize as Software Only, and IBM Storwize storage systems, and to IBM FlashSystem devices that run IBM Spectrum Virtualize.

Restriction: Key performance indicators are provided only for IBM FlashSystem 9100 from the IBM FlashSystem family. Key performance indicators are not provided for file storage in Storwize V7000 Unified storage systems.

1. To select a storage system, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems

2. From the list of storage systems that are displayed, right-click the resource that you want to view and click **View Details**.

For example, select a SAN Volume Controller or a Storwize storage system on the **Block Storage Systems** page, or select a Storwize V7000 Unified storage system on the **File Storage Systems** page.

- 3. In the **General** section of the resource details page, click **Performance**.
 - The key performance indicators charts are displayed with data that was collected from monitored resources over the last 24 hours. The performance of each monitored resource is charted against the best practice guidelines for IBM Spectrum Virtualize storage systems.
- 4. Optional: If you want to view the performance charts for a period other than the default 24-hour period, change the date for all of the charts. Click **Last 24 hours** and use the calendar to specify a date.
- 5. Optional: Drill down to view detailed information about a specific I/O group, port, or node. Click a resource from the chart's list.

For example, select the name of an I/O group in the **Read Response Time by I/O Group** chart to view the average number of milliseconds taken by the I/O group to complete a read operation.

By default, metrics are shown on the charts for the following time periods:

The last 24 hours

The same 24-hour period for yesterday

The same 24-hour period for the day seven days before today

If you selected a date from the calendar, the metrics are shown on the charts for the following time periods:

The selected date

The day before the selected date

The day seven days before the selected date

Key performance indicators

You can view the key metrics that are outside of a standard range for resources that run IBM Spectrum Virtualize by using the key performance indicators charts. The **Best Practice** values on the charts were established by a historical analysis of storage environments.

In this documentation, IBM Spectrum Virtualize is used to refer collectively to IBM SAN Volume Controller, IBM Spectrum Virtualize for Public Cloud, IBM Spectrum Virtualize as Software Only, and IBM Storwize storage systems, and to IBM FlashSystem devices that run IBM Spectrum Virtualize.

Restriction: Key performance indicators are provided only for systems that run IBM Spectrum Virtualize.

Most of the performance charts show an orange line that indicates the best practice value for the metric. These best practice values are established as the levels that allow for a diverse set of workload characteristics while maintaining a stable performance profile. The other lines on each chart represent the measured values for the metric for the resources on your storage system: I/O groups, ports, or nodes. You can use the lines to compare how close to potentially becoming overloaded your resources are. If your storage system is responding poorly and the charts indicate overloaded resources, you might have to better balance the workload. You can balance the workload between the nodes of the cluster, potentially adding more nodes to the cluster, or move some workload to other storage systems.

The key performance indicator charts show the hourly performance data measured for each resource on the selected day. Use the following charts to compare the workloads on your storage system for these key performance indicators with the best practice guidelines:

Node Utilization Percentage by Node

Compare the guideline value for this metric, for example, 60 percent utilization, with the measured value from your system.

Overall Port Bandwidth Percentage by Port

Compare the guideline value for this metric, for example, 50 percent, with the measured value from your system. Because a cluster can have many ports, the chart shows only the eight ports with the highest average bandwidth over the selected day.

Port-to-Local Node Send Response Time by Node

Compare the guideline value for this metric, for example, 0.6 ms/op, with the measured value from your system.

Port-to-Remote Node Send Response Time by Node

Because latencies for copy-services operations can vary widely, a guideline is not established for this metric. Use this chart to identify any discrepancies between the data rates of different nodes.

Read Response Time by I/O Group

Compare the guideline value for this metric, for example, 15 ms/op, with the measured value from your system.

System CPU Utilization by Node

Compare the guideline value for this metric, for example, 70 percent utilization, with the measured value from your system.

Total Data Rate by I/O Group

Because data rates can vary widely, a guideline isn't established for this metric. Use this chart to identify any significant discrepancies between the data rates of different I/O groups because these discrepancies indicate that the workload is not balanced.

Write Response Time by I/O Group

Compare the guideline value for this metric, for example, 5 ms/op, with the measured value from your system.

Zero Buffer Credit Percentage by Node

Compare the guideline value for this metric, for example, 20 percent, with the measured value from your system.

Restriction: This chart is available only for 8 Gbps Fibre Channel ports on storage systems that run IBM Spectrum Virtualize.

Port Send Delay Time

Because delay times can vary significantly depending on configuration and usage, a guideline threshold is not established for this metric. Use this chart to identify consistent discrepancies between ports and any spikes that might correlate with the time of any reported performance problems.

Available for: IBM FlashSystem devices that run IBM Spectrum Virtualize only.

Tip: The guidelines are not thresholds, and they are not related to the alerting feature in IBM Spectrum Control. To create performance alerts that use the guidelines as thresholds, go to a resource detail page in the GUI, click **Alerts** in the General section, and then click **Definitions**.

Viewing performance alerts

View the alerts that were generated when the measured value of a performance metric falls outside of the specified threshold for a resource.

To generate performance alerts, you must collect performance data and define performance alerts for that resource, or for the alert policy that manages the resource. For information about how to collect performance data, see "Creating performance monitors in IBM Spectrum Control" on page 144. For information about how to define performance alerts, see "Defining alert definitions for performance changes" on page 162.

1. Complete one of the following tasks to view performance alerts:

Task	Steps
View the performance alerts for a specific storage system	 a. In the menu bar, choose one of the following options: Storage > Block Storage Systems Storage > File Storage Systems b. Double-click a storage system. c. Click Alerts in the General section.
View all the performance alerts that were generated for switches	a. In the menu bar, select Network > Switches . b. Click the Alerts tab.
View the performance alerts for a specific switch	a. In the menu bar, select Network > Switches . b. Double-click a switch. c. Click Alerts in the General section.

2. Right-click the performance alert that you want to view and select View Alert.

In this view, a chart shows the performance of the resource that violated the threshold and generated the alert. The time range of the chart is 2 hours before and 2 hours after the violation occurred. The lines on the chart represent the following values:

- The blue line represents the performance of the resource for the selected metric.
- The red line represents the threshold that was violated.
- The dotted line represents the time when the violation occurred.

For certain storage systems, a ranked list of the 25 volumes that were the top contributors to the workload on the resource where the threshold violation occurred might be displayed below the chart. This list is shown under the following conditions:

- The performance threshold was violated on an internal resource other than a port.
- The performance threshold was violated on any of the following storage systems:
 - DS6000
 - DS8000
 - XIV

- IBM Spectrum Accelerate
- ESS
- SAN Volume Controller
- Storwize V3500
- Storwize V3700
- Storwize V7000
- Storwize V7000 Unified

You can also view the servers or hosts to which those volumes are assigned.

3. Optional: To access the full performance view of a resource that violated the threshold, click the open performance view icon 🖾 on the chart.

Exporting performance data for storage systems and fabrics

To help resolve performance issues with storage systems and fabrics, you can export performance data for the resources to a compressed file. If you contact IBM Support to help you analyze the problem, you might be asked to send this file.

You can export performance data for the following types of resource: storage systems, nodes, modules, I/O groups, host connections, pools, RAID arrays, managed disks, disks, volumes, file systems, switches, ports, and inter-switch connections.

The summarization type for the exported data can be sample, hourly, or daily.

The performance data is created as CSV files, one file for each resource type. By default, the files are compressed as .zip files.

Related tasks

Tutorial: Exporting performance data for a SAN Volume Controller system

In this tutorial, you export performance data for a SAN Volume Controller to a compressed package. You then send the package to IBM Support.

Exporting performance data by using the GUI

You can export performance data for a managed resource. If you contact IBM Support to help you analyze a performance problem with storage systems or fabrics, you might be asked to send this data.

The performance data might be large, especially if the data is for storage systems that have many volumes, or the performance monitors are running with a 1-minute sampling frequency. If the time range for the data is greater than 12 hours, volume data and 1-minute sample data is automatically excluded from the performance data even if it is available. To include volume data and 1-minute sample data, select the **Advanced package** option when you export performance data.

When you export performance data, you can specify a time range to export performance data for. The time range cannot exceed the history retention limit for sample performance data. By default, this history retention limit is two weeks.

To export hourly or daily performance data, use the exportPerformanceData script. However, the time range still cannot exceed the history retention limits for the type of performance data.

In the menu bar, select the type of storage system.
 For example, to create a compressed file for a block storage system, go to Storage > Block Storage Systems.

To create a compressed file for a fabric, go to **Network > Fabrics**.

2. Right-click the storage resource, and then click **Export Performance data**.

Exporting performance data by using a script

Use the exportPerformanceData script to export performance data for managed resources. If you contact IBM Support to help you analyze a performance problem with storage systems or fabrics, you might be asked to send this data.

Running the exportPerformanceData script

You can run the exportPerformanceData script on the server where IBM Spectrum Control is installed. The script is run from the *installation_directory*/scripts directory on the server. You can export the performance data for one or more managed resources.

The exportPerformanceData script can be run on Windows and AIX/Linux servers:

- Windows default location: C:\Program Files\IBM\TPC\scripts \exportPerformanceData.bat
- AIX/Linux default location: /opt/IBM/TPC/scripts/exportPerformanceData.sh or /usr/IBM/TPC/scripts/exportPerformanceData.sh

The following arguments are required when you run the **exportPerformanceData** script:

-user user name

Specifies an IBM Spectrum Control user ID.

-pwd password

Specifies the password for the IBM Spectrum Control user ID.

-resNames resource_name1 resource_name2 resource_name3...

Specifies the names of storage system, switch, or fabric names as displayed in the UI, separated by spaces. If the resources have spaces in the names, then you must enclose the resource names in double quotation marks (").

To view the optional parameters that the script uses, run the script with the -help option.

Examples

To collect and export performance data to a compressed file for a storage system:

```
>cd c:\Program Files\IBM\TPC\scripts\
>exportPerformanceData -user scUser1 -pwd scPwd1
   -out "c:\web data\performance reports\Array1.20160402"
   -resTypes storageSystem,Nodes,Pools,Ports
   -start "2016-02-04 00:00:00" -end "2016-02-04 23:59:59"
   -summType hourly -advPkg yes -zip no
   -resNames "Production Array 1"
```

This command creates a directory Array1.20160402 in c:\web data\performance reports\ if the directory does not exist already. The command writes to that directory the CSV files that contain the hourly performance data for the whole storage system, nodes, pools, and ports. The data is for the storage system named Production Array 1. The time range of the data is from Apr 2, 2016 0:00 AM to Apr 2, 2016 11:59 PM. One minute sample and volume data are included where they are available in the specified time range.

Performance metrics

IBM Spectrum Control can collect information about the performance of storage systems and switches. This information includes metrics that measure the performance of volumes, ports, and disks. You can view performance information or specify alerts for when the performance of a resource falls outside a specified range.

You can view performance metrics at the following locations:

Resource pages

To access a view of performance metrics for a resource, go to the following pages:

- Storage systems page, Storage system details page
- Volumes page

- · Pools page
- Switches page, Switch details page
- Server details page > Related Resources
- Hypervisor details page > Related Resources

Cognos Analytics reports

To view detailed performance reports for multiple resources in the Cognos Analytics reporting tool, go to the following URL:

http://hostname:port/bi

The *hostname* is the server that Cognos Analytics is running on. The default value for *port* is 9300. For example, to view performance reports, go to the following URL:

http://myhostname:9300/bi

Performance metrics for DS6000 and DS8000

Monitor the performance metrics that are collected for DS8000 and DS6000 storage systems.

Overview

Performance metrics are available for the following resources:

- "Volume metrics" on page 278
- "Disk metrics" on page 282
- "Port metrics" on page 284

A performance metric might apply to one or more storage resources. To check which resources the performance metric applies to, see the tips and the table footnotes.

Volume metrics

Volume performance metrics are divided into the following categories:

- Key volume metrics
- I/O rate metrics
- · Cache hit percentage metrics
- Average transfer size and HPF I/O metrics

Tip: Unless otherwise noted, you can view the volume metrics in <u>Table 120 on page 278</u>, <u>Table 121 on page 279</u>, <u>Table 122 on page 281</u>, and <u>Table 123 on page 282</u> for the following resources:

- · Host connections
- Nodes
- Pools
- · Storage systems
- Volumes

Table 120. Key volume metrics	
Metric	Description
Cache Holding Time <u>"1" on page 279</u>	The average number of seconds that I/O data for a storage system node is held in the cache. A short cache-holding time indicates adverse performance.

Table 120. Key volume metrics (continued, Metric	Description
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Pool Activity Score "2" on page 279	The activity level of pools, which is set to the following value:
	<pre>[Read I/O Rate × (1- Read I/O Cache Hit %)] ÷ Total Pool Capacity</pre>
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
Volume Utilization "3" on page 279	The average percentage of time that the volume is busy.
Write-Cache Delay Percentage	The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions. The value is a percentage of all operations.

Notes:

- 1. This metric is only available when you view the performance of nodes or storage systems.
- 2. This metric is only available when you view the performance of pools.
- 3. This metric is only available when you view the performance of volumes.

Table 121. I/O Rates	
Metric	Description
Average Transfer Rate (Cache-to-Disk) 1" on page 280	The average number of sectors or tracks per second that are transferred from the cache to the disks.

Table 121. I/O Rates (continued)	
Description	
The percentage of cache hits for record-mode read operations. For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.	
The average number of read operations per second that are issued by the High Performance FICON feature of the storage system.	
The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.	
The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.	
The average number of nonsequential read operations per second.	
The average number of nonsequential write operations per second.	
The average number of nonsequential I/O operations per second. This value includes both read and write operations.	
The average number of tracks per second that are transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.	
The average number of I/O operations per second for record-mode read operations. For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.	
The average number of sequential read operations per second.	
The average number of sequential write operations per second.	
The average number of sequential I/O operations per second. This value includes both read and write operations.	
The average number of I/O operations per second that are delayed because of space constraints in the write cache, or because of other conditions.	

Notes:

- 1. This metric is only available when you view the performance of volumes, pools, nodes, and host connections.
- 2. This metric applies only to DS8000 storage systems.

Table 122. Cache hit percentages Metric Description	
Normal I/O Cache Hits (Read)	The percentage of nonsequential read operations that find data in the cache. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Normal I/O Cache Hits (Write)	The percentage of nonsequential write operations that are handled in the cache.
Normal I/O Cache Hits (Total)	The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Overall I/O Cache Hits (Read)	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.
Overall I/O Cache Hits (Write)	The average percentage of all write operations that are handled in the cache, across all volumes on the server. This value includes both sequential and nonsequential write operations.
Overall I/O Cache Hits (Total)	The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Record Mode Read Cache Hit Percentage <u>"1" on page 282</u>	The percentage of cache hits for record-mode read operations. For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.
Sequential I/O Cache Hits (Read)	The percentage of sequential read operations that find data in the cache. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Table 122. Cache hit percentages (continued)	
Metric	Description
Sequential I/O Cache Hits (Write)	The percentage of sequential write operations that are handled in the cache.
Sequential I/O Cache Hits (Total)	The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Note:

1. This metric is only available when you view the performance of volumes, pools, nodes, and host connections.

Table 123. Average transfer size, HPF I/O, and miscellaneous metrics	
Metric	Description
Average Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Average Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Average Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.
HPF I/O Percentage	The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.

Disk metrics

Disk performance metrics are divided into the following categories:

- Key disk metrics
- Transfer size metrics

Tip: Unless otherwise noted, you can view the disk metrics in <u>Table 124 on page 282</u> and <u>Table 125 on page 284</u> for the following resources:

- Nodes
- Pools
- RAID arrays
- Ranks
- · Storage systems

Table 124. Key disk metrics	
Metric	Description
Data Rate (Read)	The average number of MiB per second that are read from the back-end storage resources.

Table 124. Key disk metrics (continued)	
Metric	Description
Data Rate (Write)	The average number of MiB per second that are written to the back-end storage resources.
Data Rate (Total)	The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.
Disk Utilization Percentage ^{"1" on page 284}	The average percentage of time that the disks that are associated with an array are busy. No value is calculated for this property if there are multiple ranks in the extent pool where the thin-provisioned volumes are allocated. In this case, the value N/A is displayed. This limitation applies only to DS8000 storage systems. If there is only a single rank in the extent pool, the value for this property is calculated regardless of the thin-provisioned volumes.
	Tip: Some highly sequential workloads such as batch or backup processing might continually exceed the threshold because they drive the arrays to high utilization percentages. For these types of workloads, a high utilization indicates that the work is being performed very efficiently and is not a cause for concern.
	For DS8000, this metric is available only for version 8.5.0 and later. For earlier versions, the value N/A is shown in the performance table view and the metric is not shown on the performance chart.
I/O Rate (Read)	The average number of read operations per second that are issued to the back-end storage resources.
I/O Rate (Write)	The average number of write operations per second that are issued to the back-end storage resources.
I/O Rate (Total)	The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.
Response Time (Read)	The average number of milliseconds for the backend storage resources to respond to a read operation.
Response Time (Write)	The average number of milliseconds for the backend storage resources to respond to a write operation.
Response Time (Overall)	The average number of milliseconds for the backend storage resources to respond to a read or a write operation.

Table 124. Key disk metrics (continued)		
Metric Description		
Note:		
1. This metric is only available when you view the performance of RAID arrays.		

Table 125. Transfer size metrics	
Metric	Description
Average Transfer Size (Read)	The average number of KiB that are transferred per read operation from the back-end storage resources.
Average Transfer Size (Write)	The average number of KiB that are transferred per write operation to the back-end storage resources.
Average Transfer Size (Overall)	The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Port metrics

Port performance metrics are divided into the following categories:

- · Key port metrics
- Port I/O rate metrics
- Port data rate metrics
- · Port response time metrics
- Error rate metrics
- · Remote mirror metrics
- Transfer size metrics

Tip: Unless otherwise noted, you can view the port metrics in <u>Table 126 on page 284</u>, <u>Table 130 on page 286</u>, <u>Table 131 on page 288</u>, and <u>Table 132 on page 289</u> for ports and storage systems. The port metrics in Table 127 on page 285, Table 128 on page 286, Table 129 on page 286 are only available for ports.

Table 126. Key port metrics	
Metric	Description
Data Rate (Receive)	The average rate at which data is received by the port. The rate is measured in MiB per second.
Data Rate (Send)	The average rate at which data is sent through the port. The rate is measured in MiB per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
I/O Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.
I/O Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.
I/O Rate (Total)	The average number of send operations and receive operations per second.

Table 126. Key port metrics (continued)	
Metric	Description
Response Time (Receive)	The average number of milliseconds to complete a receive operation.
Response Time (Send)	The average number of milliseconds to complete a send operation.
Response Time (Overall)	The average number of milliseconds to complete a send or receive operation.
Port Utilization (Receive) 10 on page 285	The average percentage of time that the port is busy receiving data.
Port Utilization (Send) "1" on page 285	The average percentage of time that the port is busy sending data.
Port Utilization (Overall) "1" on page 285	The average percentage of time that the port is busy sending or receiving data.
Bandwidth (Receive) "1" on page 285	The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth (Send) ^{"1" on page 285}	The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth (Overall) ^{"1" on page 285}	The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Note:

1. This metric is only available when you view the performance of ports.

Table 127. Port I/O rate metrics	
Metric	Description
FICON I/O Rate (Send)	The average number of send operations per second for Fibre Channel connection (FICON) usage.
FICON I/O Rate (Receive)	The average number of receive operations per second for FICON usage.
FICON I/O Rate (Total)	The average number of send and receive operations per second for FICON usage.
FCP I/O Rate (Send)	The average number of send operations per second for Fibre Channel Protocol (FCP) usage.
FCP I/O Rate (Receive)	The average number of receive operations per second for FCP usage.
FCP I/O Rate (Total)	The average number of send operations and receive operations per second for FCP usage.

Table 128. Port data rates	
Metric	Description
FICON Data Rate (Send)	The average number of MiB per second that is sent for FICON usage.
FICON Data Rate (Receive)	The average number of MiB per second that is received for FICON usage.
FICON Data Rate (Total)	The average number of MiB per second that is transferred for FICON usage. This value includes both send and receive FICON operations.
FCP Data Rate (Send)	The average number of MiB per second that are sent for FCP usage.
FCP Data Rate (Receive)	The average number of MiB per second that are received for FCP usage.
FCP Data Rate (Total)	The average number of MiB per second that are transferred for FCP usage. This value includes both send and receive FCP operations.

Table 129. Port response times	
Metric	Description
FICON Response Time (Send)	The average number of milliseconds to complete a send operation for FICON usage.
FICON Response Time (Receive)	The average number of milliseconds to complete a receive operation for FICON usage.
FICON Response Time (Overall)	The average number of milliseconds to complete a send or receive operation for FICON usage. This value includes both send and receive FICON operations.
FCP Response Time (Send)	The average number of milliseconds to complete a send operation for FCP usage.
FCP Response Time (Receive)	The average number of milliseconds to complete a receive operation for FCP usage.
FCP Response Time (Overall)	The average number of milliseconds to complete a send or receive operation for FCP usage.

Table 130. Error rate metrics.	
Tip: You can view the following metrics for ports and storage systems.	
Metric	Description
Frame Errors (CRC Errors)	The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver. For Brocade switches, this metric includes only the CRC Errors with a good end-of-frame (EOF) indicator.

Table 130. Error rate metrics.

Tip: You can view the following metrics for ports and storage systems.

(continued)

(continuea)	
Metric	Description
Frame Errors (Error Frame)	The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.
Frame Errors (Invalid Relative Offset Rate)	The average number of times per second that frames are received with an invalid relative offset in the frame header.
Link Errors (Invalid Link Transmission)	The average number of times per second that an invalid transmission word was detected by the port while the link did not experience any signal or synchronization loss.
Link Errors (Invalid Transmission Words)	The average number of bit errors per second that are detected.
Link Errors (Link Failures)	The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.
Link Errors (Primitive Sequence Protocol Error Rate)	The average number of primitive sequence protocol errors per second that are detected. This error occurs when there is a link failure for a port.
Link Errors (Sequence Timeouts)	The average number of times per second that the port detects a timeout condition after the port receives a sequence initiative for a Fibre Channel exchange.
Link Errors (Signal Loss)	The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.
Link Errors (Sync Loss)	The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.

Table	p 130	Frror rate	metrics

Tip: You can view the following metrics for ports and storage systems.

(continued)

Metric	Description
Port Protocol Errors (Credit Recovery Link Resets)	The estimated average number of link resets per second that a switch or port completed to recover buffer credits. This estimate attempts to disregard link resets that were caused by link initialization. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.
Port Protocol Errors (Duplicate Frames)	The average number of duplicate frames per second that are received. A duplicate frame is a frame that the system previously processed for the port.
Port Protocol Errors (Link Reset Received)	The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state.
Port Protocol Errors (Link Reset Transmitted)	The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state.
Port Protocol Errors (Out of Order ACK)	The average number of times per second that an out-of-order acknowledge (ACK) frame is detected. An ACK frame is used for end-to-end flow control and is sent to verify receipt of a frame.
Port Protocol Errors (Out of Order Data)	The average number of times per second that an out-of-order frame is detected.

Note:

1. This metric is only available when you view the performance of ports.

Table 131. Remote mirror metrics	
Metric	Description
PPRC Data Rate (Receive)	The average number of MiB per second that are received by using the Peer-to-Peer Remote Copy protocol.
PPRC Data Rate (Send)	The average number of MiB per second that are sent by using the Peer-to-Peer Remote Copy protocol.
PPRC Data Rate (Total)	The average number of MiB per second that are transferred by using the Peer-to-Peer Remote Copy (PPRC) protocol. This value includes both send and receive PPRC operations.
PPRC I/O Rate (Receive)	The average number of operations per second that are received by using the Peer-to-Peer Remote Copy protocol.

Table 131. Remote mirror metrics (continued)	
Metric	Description
PPRC I/O Rate (Send)	The average number of operations per second that are sent by using the Peer-to-Peer Remote Copy protocol.
PPRC I/O Rate (Total)	The average number of send operations and receive operations per second using the Peer-to-Peer Remote Copy protocol.
PPRC Response Time (Receive)	The average number of milliseconds to complete a receive operation by using the Peer-to-Peer Remote Copy protocol.
PPRC Response Time (Send)	The average number of milliseconds to complete a send operation by using the Peer-to-Peer Remote Copy protocol.
PPRC Response Time (Overall)	The average number of milliseconds to complete a send or receive operation by using the Peer-to-Peer Remote Copy (PPRC) protocol.

Table 132. Transfer size metrics	
Metric	Description
Average Transfer Size (Receive)	The average number of KiB that are transferred per receive operation.
Average Transfer Size (Send)	The average number of KiB that are transferred per send operation.
Average Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Performance metrics for resources that run IBM Spectrum Virtualize

Monitor the performance metrics that are collected for IBM Spectrum Virtualize storage systems.

Overview

In this documentation, IBM Spectrum Virtualize is used to refer collectively to IBM SAN Volume Controller, IBM Spectrum Virtualize for Public Cloud, IBM Spectrum Virtualize as Software Only, and IBM Storwize storage systems, and to IBM FlashSystem devices that run IBM Spectrum Virtualize.

Definitions are provided for the performance metrics that are collected for the following storage systems:

- FlashSystem 5000
- FlashSystem 5100
- FlashSystem 7200
- FlashSystem 9100
- FlashSystem 9200
- FlashSystem V840
- FlashSystem V9000
- SAN Volume Controller
- IBM Spectrum Virtualize for Public Cloud
- Storwize V3500

- Storwize V3700
- Storwize V5000
- Storwize V7000
- Storwize V7000 Unified (block storage only)

The following terms are used in the performance metrics for these storage systems:

Stage

To write data from a disk to the cache. The data is not prefetched data.

Destage

To write data from the cache to a disk.

Prestage

To write prefetched data from a disk to the cache.

The performance metrics are divided into the following categories:

- "Volume performance metrics" on page 290
- "Disk performance metrics" on page 300
- "Pool performance metrics" on page 302
- "Port performance metrics" on page 303
- "Node performance metrics" on page 309

Volume performance metrics

Volume performance metrics are divided into the following categories:

- Key metrics for volumes
- I/O rate metrics for volumes
- Cache hit percentage metrics for volumes
- Response time metrics for volumes
- Remote mirror metrics for volumes
- Volume cache (VC) metrics for volumes
- Volume copy cache (VCC) metrics for volumes
- Compression metrics for volumes
- Miscellaneous metrics for volumes
- Legacy cache metrics for volumes

Tip:

Unless otherwise noted, you can view the volume metrics in <u>Table 133 on page 291</u>, <u>Table 134 on page 292</u>, <u>Table 135 on page 293</u>, <u>Table 136 on page 293</u>, <u>Table 137 on page 293</u>, <u>Table 138 on page 294</u>, <u>Table 139 on page 296</u>, <u>Table 141 on page 298</u>, and <u>Table 142 on page 299</u> for the following resources:

- · Host connections
- I/O groups
- Nodes
- Pools
- · Storage systems
- Volumes

Table 133. Key metrics for volumes		
Metric	Definition	
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.	
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.	
Data Rate (Unmap) <u>"1" on page 292</u>	The average number of MiBs per second that were unmapped. This metric corresponds to the collected ub statistic.	
Data Rate (Total)	The average number of MiB per second that are transferred for read, write, and unmap operations. "1" on page 292	
Overall Host Attributed Response Time Percentage	The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently. For read response time, the value is based on the time that it takes for hosts to respond to transfer-ready notifications from the nodes. For write response time, the value is based on the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.	
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.	
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.	
Overall I/O Rate (Unmap) ^{"1" on page} 292	The average number of unmap operations per second. This metric corresponds to the collected uo statistic.	
Overall I/O Rate (Total)	The average number of nonsequential I/O operations per second. This value includes read, write, and unmap operations. "1" on page 292	
Pool Activity Score ²	The activity level of pools, which is set to the following value: [Read I/O Rate × (1- Read I/O Cache Hit %)] ÷ Total Pool Capacity	
Response Time (Read)	The average number of milliseconds to complete a read operation.	
Response Time (Write)	The average number of milliseconds to complete a write operation.	
Response Time (Unmap) ^{"1" on page 292}	The average number of milliseconds required to complete an unmap operation. This metric corresponds to the collected ul statistic.	
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes read, write, and unmap operations. "1" on page 292	
Volume Utilization ³	The average percentage of time that the volume is busy.	

Table 133. Key metrics for volumes (continued)	
Metric	Definition
Write Cache Delay	The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions. The value is a percentage of all operations. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than V7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are V7.3 or later with earlier versions.

- 1. This metric applies only to storage systems that are running IBM Spectrum Virtualize V8.1.1 or later. To view details about collected statistics, see Starting statistics collection.
- 2. This metric is also available when you view the performance of pools.
- 3. This metric is only available when you view the performance of volumes.

Table 134. I/O rate metrics for volumes	
Metric	Definition
Transfer Rate (Cache-to-Disk)*	The average number of sectors or tracks per second that are transferred from the cache to the disks.
Transfer Rate (Disk-to-Cache)*	The average number of sectors or tracks per second that are transferred per second from the disks to the cache.
Unaligned Unmap I/O Rate	The average number of volume unmap operations per second that are not aligned on an 8K boundary. This metric corresponds to the collected uou statistic. This metric applies only to storage systems that are running IBM Spectrum Virtualize V8.1.1 or later. To view details about collected statistics, see Starting statistics collection.
Write-Cache Delay I/O Rate*	The average number of I/O operations per second that are delayed because of space constraints in the write cache, or because of other conditions. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than V7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are V7.3 or later with earlier versions.

Note: *This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than V7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are V7.3 or later with earlier versions.

Table 135. Cache hit percentages metrics for volumes	
Metric	Definition
Overall I/O Cache Hits (Read)	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.
Overall I/O Cache Hits (Write)	The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.
Overall I/O Cache Hits (Total)	The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Table 136. Response time metrics for volumes	
Metric	Definition
Peak Response Time (Read)	The worst response time measured for a read operation in the sample interval.
Peak Response Time (Write)	The worst response time measured for a write operation in the sample interval.
Peak Response Time (Unmap)	The worst response time measured for an unmap operation in the sample interval. This metric corresponds to the collected ulw statistic. This metric applies only to storage systems that are running IBM Spectrum Virtualize V8.1.1 or later. To view details about collected statistics, see Starting statistics collection.

Table 137. Remote mirror metrics for volumes	
Metric	Definition
Global Mirror (Overlapping Write I/O Rate)	The average number of overlapping write operations per second that are issued by the Global Mirror primary site. Some overlapping writes are processed in parallel and are excluded from this value.

Table 137. Remote mirror metrics for volumes (continued)	
Metric	Definition
Global Mirror (Overlapping Write Percentage)	The percentage of overlapping write operations that are issued by the Global Mirror primary site. Some overlapping writes are processed in parallel and are excluded from this value.
Global Mirror (Secondary Write Lag)	The average number of additional milliseconds that it takes to service each secondary write operation for Global Mirror. This value does not include the time to service the primary write operations. You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.
Global Mirror (Write I/O Rate)	The average number of write operations per second that are issued to the Global Mirror secondary site.

Volume cache (VC) metrics are only available for SAN Volume Controller, Storwize, and FlashSystem block storage systems whose firmware version is V7.3 or later.

Tip:

The volume cache is sometimes referred to as *upper cache*.

Table 138. Volume cache (VC) metrics for volumes	
Metric	Definition
Cache Hits (Dirty Writes)	The percentage of all cache write hits that occur on data that is marked as modified in the volume cache. This value represents how effectively write operations are coalesced before the data is written to disk.
Cache Hits (Read)	The percentage of read operations that find data in the volume cache.
Cache Hits (Total)	The percentage of cache hits for read and write operations that are handled in the volume cache.
Cache Hits (Write)	The percentage of cache hits for write operations that are handled in the volume cache.
I/O Rate (Destage)	The average number of cache-to-disk transfer operations per second that are processed in the volume cache.
I/O Rate (Stage)	The average number of disk-to-cache transfer operations per second that are processed in the volume cache.
I/O Rate (Read)	The average number of read operations per second that are processed in the volume cache. This value includes operations that are started by hosts or by remote replication sources.

Table 138. Volume cache (VC) metrics for volumes (continued) Metric Definition	
I/O Rate (Total)	The average number of read and write operations per second that are processed by the volume cache. This value includes requests sent by hosts and by remote replication sources.
I/O Rate (Write)	The average number of write operations per second that are processed by the volume cache. This value includes operations that are started by hosts or by remote replication sources.
Response Time (Destage)	The average number of milliseconds that it took to complete each destage operation in the volume cache. That is, the time that it took to do write operations from the volume cache to the disk.
Response Time (Stage)	The average number of milliseconds that it took to complete each stage operation in the volume cache. That is, the time that it took to do read operations from the disk to the volume cache.
Transfer Rates (Cache-to-Disk)	The average number of sectors that are transferred per second from the volume cache to the disks.
Transfer Rates (Disk-to-Cache)	The average number of sectors that are transferred per second from the disks to the volume cache.
Write Delay Percentage (Flush-through)	The percentage of write operations that are written to disk in flush-through mode in the volume cache.
Write Delay Percentage (Total Delay)	The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions in the volume cache. The value is a percentage of all operations.
Write Delay Percentage (Write-through)	The percentage of write operations that are written to disk in write-through mode in the volume cache.
Write Delay Rate (Flush-through)	The average number of tracks per second that are written to disk in flush-through mode in the volume cache.
Write Delay Rate (Total Delay)	The average number of I/O operations per second that are delayed. The delay might occur because of space constraints in the write cache, or because of other conditions in the volume cache. The value is an average of all operations.
Write Delay Rate (Write-through)	The average number of sectors per second that are written to disk in write-through mode in the volume cache.

Volume copy cache (VCC) metrics are only available for SAN Volume Controller, Storwize, and FlashSystem block storage systems whose firmware version is V7.3 or later.

Tip: The volume copy cache is sometimes referred to as *lower cache*.

Table 139. Volume copy cache (VCC) metrics for volumes	
Metric	Definition
Cache Hits (Dirty Writes)	The percentage of all cache write hits that occur on data that is marked as modified in the volume copy cache. This value represents how effectively write operations are coalesced before the data is written to disk.
Cache Hits (Read-ahead)*	The percentage of all read cache hits that occur on pre-staged data.
Cache Hits (Read)	The percentage of read operations that find data in the volume copy cache.
Cache Hits (Total)	The percentage of cache hits for read and write operations that are handled in the volume copy cache.
Cache Hits (Write)	The percentage of cache hits for write operations that are handled in the volume copy cache.
I/O Rate (Destage)	The average number of cache-to-disk transfer operations per second that are processed in the volume copy cache.
I/O Rate (Prestage)	The average number of prefetch disk-to-cache transfer operations per second that are processed in the volume copy cache.
I/O Rate (Stage)	The average number of disk-to-cache transfer operations per second that are not prefetch operations and are processed in the volume copy cache.
I/O Rate (Read)	The average number of read operations per second that are processed in the volume copy cache. This value includes read operations that are associated with FlashCopy® services, volume mirroring, and other internal processes. This value might also include some operations that are passed from the volume cache.
I/O Rate (Total)	The average number of read and write operations per second that are processed by the volume copy cache. This value includes operations that are associated with FlashCopy services, volume mirroring, and other internal processes. This value might also include some operations that are passed from the volume cache.
I/O Rate (Write)	The average number of write operations per second that are processed by the volume copy cache. This value includes read operations that are associated with FlashCopy services, volume mirroring, and other internal processes. This value might also include some operations that are passed from the volume cache.

Metric	Definition
Response Time (Destage)	The average number of milliseconds that it took to complete each destage operation in the volume copy cache. That is, the time that it took to do write operations from the volume copy cache to the disk.
Response Time (Prestage)	The average number of milliseconds that it took to complete each prestage operation in the volume copy cache. That is, the time that it took to prefetch data from the disk into the volume copy cache.
Response Time (Stage)	The average number of milliseconds that it took to complete each stage operation in the volume copy cache. That is, the time that it took to do read operations from the disk to the volume copy cache.
Transfer Rates (Cache-to-Disk)	The average number of sectors that are transferred per second from the volume copy cache to the disks.
Transfer Rates (Disk-to-Cache)	The average number of sectors that are transferred per second from the disks to the volume copy cache.
Write Delay Percentage (Flush-through)	The percentage of write operations that are written to disk in flush-through mode in the volume copy cache.
Write Delay Percentage (Total Delay)	The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions in the volume copy cache. The value is a percentage of all operations.
Write Delay Percentage (Write-through)	The percentage of write operations that are written to disk in write-through mode in the volume copy cache.
Write Delay Rate (Flush-through)	The average number of sectors per second that are written to disk in flush-through mode in the volume copy cache.
Write Delay Rate (Total Delay)	The average number of I/O operations per second that are delayed. The delay might occur because of space constraints in the write cache, or because of other conditions in the volume copy cache. The value is an average of all operations.
Write Delay Rate (Write-through)	The average number of sectors per second that are written to disk in write-through mode in the volume copy cache.

Note: *This metric is only available for SAN Volume Controller, Storwize, and FlashSystem block storage systems whose firmware version is V7.4 or later.

Note:

Unless otherwise noted, you can view the volume metrics in <u>Table 140 on page 298</u> for the following resources:

- Nodes
- I/O groups
- Storage systems

Table 140. Compression metrics for volumes	
Metric	Definition
Compressed Volumes I/O Rate	The average number of all read and write operations per second for compressed volumes.
Compressed Volumes Data Rate	The average number of MiB per second that were read from or written to compressed volumes.
Compressed Volumes Response Time	The average number of milliseconds to complete an I/O operation for compressed volumes. This value includes both read and write operations.
Uncompressed Volumes I/O Rate	The average number of all read and write operations per second for uncompressed volumes.
Uncompressed Volumes Data Rate	The average number of MiB per second that were read from or written to uncompressed volumes.
Uncompressed Volumes Response Time	The average number of milliseconds to complete an I/O operation for uncompressed volumes. This value includes both read and write operations.

Tip:

Unless otherwise noted, you can view the volume metrics in <u>Table 141 on page 298</u> for the following resources:

- Nodes
- I/O groups
- Host connections
- Storage systems

Table 141. Miscellaneous metrics for volumes	
Metric	Definition
Cache to Host Transfer Response Time "1" on page 299	The average number of milliseconds that is taken to transfer a track from the cache to the host, including any queuing time that occurs because of throttling.
Non-Preferred Node Usage Percentage "2" on page 299	The overall percentage of I/O operations that are not directed against the preferred node for each volume in an I/O Group. There is a small performance penalty when I/O does not go to the preferred node for each volume.
Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Table 141. Miscellaneous metrics for volumes (continued)	
Metric	Definition
Unaligned Write I/O Rate ^{"3" on page 299}	The average number of write operations per second that are not aligned on a boundary between 4 KB physical blocks. Unaligned write operations can cause a significant decrease in efficiency of write operations to some types of back-end disks. You can ignore this metric for iSeries servers.

- 1. The metric is only available for SAN Volume Controller, Storwize, and FlashSystem block storage systems whose firmware version is V7.3 or later.
- 2. This metric is only available when you view the performance of volumes, I/O groups, and host connections.
- 3. This metric is only available when you view the performance of pools, volumes, and host connections. The metric is only available for SAN Volume Controller, Storwize, and FlashSystem block storage systems whose firmware version is V7.4 or later.

Legacy cache metrics are only available for SAN Volume Controller and Storwize block storage systems whose firmware version is earlier than V7.3.

Table 142. Legacy cache metrics for volumes	
Metric	Definition
Write-Cache Percentages (Flush-through)	The percentage of tracks that are written to disk in flush-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than V7.3.
Write-Cache Percentages (Write-through)	The percentage of tracks that are written to disk in write-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than V7.3.
Write-Cache Percentages (Overflow)	The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.
Write-Cache I/O Rate (Flush-through)	The average number of tracks per second that are written to disk in flush-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than V7.3.
Write-Cache I/O Rate (Overflow)	The average number of tracks per second that are written but are delayed because there is not enough space in the write cache. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than V7.3.
Write-Cache I/O Rate (Write-through)	The average number of tracks per second that are written to disk in write-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than V7.3.

Table 142. Legacy cache metrics for volumes (continued)	
Metric	Definition
Dirty Write Percentage of Cache Hits	The percentage of all cache write hits that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than V7.3.

Disk performance metrics

Disk performance metrics are divided into the following categories:

- · Key metrics for disks
- Response time metrics for disks
- Miscellaneous metrics for disks

Unless otherwise noted, you can view disk metrics for the following resources:

- Managed disks
- Pools
- Nodes
- I/O Groups
- Storage systems

Restriction: Performance metadata for managed disks in IBM Spectrum Virtualize for Public Cloud is not yet available.

Table 143. Key metrics for disks	
Metric	Definition
Data Rate (Read)	The average number of MiB per second that are read from the back-end storage resources.
Data Rate (Write)	The average number of MiB per second that are written to the back-end storage resources.
Data Rate (Total)	The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.
I/O Rate (Read) ¹	The average number of read operations per second that are issued to the back-end storage resources.
I/O Rate (Write) ²	The average number of write operations per second that are issued to the back-end storage resources.
I/O Rate (Total) ³	The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Table 143. Key metrics for disks (continued)	
Metric	Definition
Response Time (Read)	The average number of milliseconds for the backend storage resources to respond to a read operation.
Response Time (Write)	The average number of milliseconds for the backend storage resources to respond to a write operation.
Response Time (Overall)	The average number of milliseconds for the backend storage resources to respond to a read or a write operation.

- 1. The performance metrics for I/O Rate (Read) are available for pools, nodes, I/O groups, and storage systems.
- 2. The performance metrics for I/O Rate (Write) are available for pools, nodes, I/O groups, and storage systems.
- 3. The performance metrics for I/O Rate (Total) are available for pools, nodes, I/O groups, and storage systems.

Table 144. Response time metrics for disks	
Metric	Definition
Queue Time (Read)	The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.
Queue Time (Write)	The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.
Queue Time (Overall)	The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.
Peak Back-end Queue Time (Read)	The longest time that a read operation spends in the queue before the operation is sent to the backend storage resources.
Peak Back-end Queue Time (Write)	The longest time that a write operation spends in the queue before the operation is sent to the backend storage resources.
Peak Back-end Response Time (Read)	The longest time for a back-end storage resource to respond to a read operation.
Peak Back-end Response Time (Write)	The longest time for a back-end storage resource to respond to a write operation by a node.

Table 145. Miscellaneous metrics for disks	
Metric	Definition
Transfer Size (Read)	The average number of KiB that are transferred per read operation from the back-end storage resources.
Transfer Size (Write)	The average number of KiB that are transferred per write operation to the back-end storage resources.
Transfer Size (Overall)	The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.
Cache Destage (In-Flight I/O)*	The actual number of concurrent operations that are issued from the cache to the managed disk. The number of in-flight operations cannot exceed the target, but also varies over time, depending on the type of I/O operations and on the performance of the managed disk.
Cache Destage (Target I/O)*	The maximum number of concurrent operations that can be issued from the cache to the managed disk. This target changes over time, depending on changes in cache utilization and on changes in the performance of the managed disk.

Note: *This metric is only available when you view the performance of managed disks. The metric is only available for SAN Volume Controller, Storwize, and FlashSystem block storage systems whose firmware version is V7.3 or later.

Pool performance metrics

Key performance metrics are available for pools.

Unless otherwise noted, you can view pool metrics for the following resources:

• Pools

Table 146. Key metrics for pools	
Metric	Definition
Max Write Cache Fullness*	The maximum amount of the lower cache that the write cache partitions on the nodes that manage the pool are using for write operations. If the value is 100%, one or more cache partitions on one or more pools is full. The operations that pass through the pools with full cache partitions will be queued and I/O response times will increase for the volumes in the affected pools.
Write Cache Fullness*	The average amount of the lower cache that the pools' write cache partitions on the nodes are using for write operations. Monitor average cache fullness to identify the pools that are experiencing heavy cache usage.

Note: *This cache fullness metric applies to systems that are running IBM Spectrum Virtualize V7.3 or later.

Port performance metrics

Port performance metrics are divided into the following categories:

- · Key metrics for ports
- I/O rate metrics for ports
- Data rate metrics for ports
- · Response time metrics for ports
- Error rate metrics for ports
- Miscellaneous metrics for ports

Unless otherwise noted, you can view port metrics for the following resources:

- Ports
- Nodes
- I/O Groups
- Storage systems

Table 147. Key metrics for ports	
Metric	Definition
I/O Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.
I/O Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.
I/O Rate (Total)	The average number of send operations and receive operations per second.
Data Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.
Data Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Table 148. I/O rate metrics for ports	
Metric	Definition
Port-to-Disk I/O Rate (Receive)	The average number of exchanges per second that are received from back-end storage resources.
Port-to-Disk I/O Rate (Send)	The average number of IOs per second that are sent from the storage system to the back-end storage it is virtualizing. Use this metric to help measure the rate of data that is sent to back-end storage.
Port-to-Disk I/O Rate (Total)	The average number of IOs per second that are transmitted between the storage system and the back-end storage it is virtualizing. Use this metric to help measure the rate of data that is sent to back-end storage.

Table 148. I/O rate metrics for ports (continued)	
Metric	Definition
Port-to-Host I/O Rate (Receive)	The average number of IOs per second that are received by the storage system from the hosts that are accessing its storage. Use this metric to help measure host workload against the storage system.
Port-to-Host I/O Rate (Send)	The average number of IOs per second that are sent by the storage system to the hosts that are accessing its storage. Use this metric to help measure host workload against the storage system.
Port-to-Host I/O Rate (Total)	The average number of IOs per second that are transmitted between the storage system and the hosts that are accessing its storage. Use this metric to help measure host workload against the storage system.
Port-to-Local Node I/O Rate (Receive)	The average number of IOs per second that are received from other nodes within the local cluster. Use this metric to understand the rate of intercluster communication.
Port-to-Local Node I/O Rate (Send)	The average number of IOs per second that are sent to other nodes within the local cluster. Use this metric to understand the rate of inter-cluster communication.
Port-to-Local Node I/O Rate (Total)	The average number of IOs per second that are transmitted between the resource and other nodes within the local cluster. Use this metric to understand the rate of inter-cluster communication.
Port-to-Remote Node I/O Rate (Receive)	The average number of IOs per second that are received from nodes that are in a remote cluster. Use this metric to understand the amount of remote replication workload.
Port-to-Remote Node I/O Rate (Send)	The average number of IOs per second that are sent to nodes that are in a remote cluster. Use this metric to understand the amount of remote replication workload.
Port-to-Remote Node I/O Rate (Total)	The average number of IOs per second that are transmitted between the resource and nodes that are in a remote cluster. Use this metric to understand the amount of remote replication workload.

Table 149. Data rate metrics for ports	
Metric	Definition
Port-to-Disk Data Rate (Receive)	The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

Table 149. Data rate metrics for ports (continued)	
Metric	Definition
Port-to-Disk Data Rate (Send)	The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.
Port-to-Disk Data Rate (Total)	The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.
Port-to-Host Data Rate (Receive)	The average rate at which data is received from host computers. The rate is measured in MiB per second.
Port-to-Host Data Rate (Send)	The average rate at which data is sent to host computers. The rate is measured in MiB per second.
Port-to-Host Data Rate (Total)	The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.
Port-to-Local Node Data Rate (Receive)	The average rate at which data is received from other nodes that are in the local cluster. The rate is measured in MiB per second.
Port-to-Local Node Data Rate (Send)	The average rate at which data is sent to other nodes that are in the local cluster. The rate is measured in MiB per second.
Port-to-Local Node Data Rate (Total)	The average rate at which data is transmitted between the component and other nodes that are in the local cluster. The rate is measured in MiB per second.
Port-to-Remote Node Data Rate (Receive)	The average rate at which data is received from nodes that are in the remote cluster. The rate is measured in MiB per second.
Port-to-Remote Node Data Rate (Send)	The average rate at which data is sent to nodes that are in the remote cluster. The rate is measured in MiB per second.
Port-to-Remote Node Data Rate (Total)	The average rate at which data is transmitted between the component and nodes that are in the remote cluster. The rate is measured in MiB per second.

Table 150. Response time metrics for ports	
Metric	Definition
Port-to-Local Node Response Time (Receive)	The average number of milliseconds to complete a receive operation from another node that is in the local cluster. This value represents the external response time of the transfers.

Table 150. Response time metrics for ports (continued)	
Metric	Definition
Port-to-Local Node Response Time (Send)	The average number of milliseconds to complete a send operation to another node that is in the local cluster. This value represents the external response time of the transfers.
Port-to-Local Node Response Time (Overall)	The average number of milliseconds to complete a send or receive operation with another node that is in the local cluster. This value represents the external response time of the transfers.
Port-to-Remote Node Response Time (Overall)	The average number of milliseconds to complete a send operation to, or a receive operation from a node in the remote cluster. This value represents the external response time of the transfers.
Port-to-Remote Node Response Time (Receive)	The average number of milliseconds to complete a receive operation from a node that is in the remote cluster. This value represents the external response time of the transfers.
Port-to-Remote Node Response Time (Send)	The average number of milliseconds to complete a send operation to a node that is in the remote cluster. This value represents the external response time of the transfers.
Port-to-Remote Node Response Time (Overall)	The average number of milliseconds to complete a send operation to, or a receive operation from a node in the remote cluster. This value represents the external response time of the transfers.

 $\begin{tabular}{ll} \textbf{Metrics availability restrictions:} The response time metrics are available for nodes, I/O groups, and storage systems. \end{tabular}$

Table 151. Error rate metrics for ports	
Metric	Definition
CRC Error Rate	The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver. For Brocade switches, this metric includes only the CRC Errors with a good end-of-frame (EOF) indicator.
Link Errors (Invalid Link Transmission Rate)	The average number of bit errors per second that are detected.
Link Errors (Invalid Transmission Word Rate)	The average number of bit errors per second that are detected.
Link Errors (Link Failures)	The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

Table 151. Error rate metrics for ports (continued)	
Metric	Definition
Link Errors (Primitive Sequence Protocol Error Rate)	The average number of primitive sequence protocol errors per second that are detected. This error occurs when there is a link failure for a port.
Link Errors (Signal Loss)	The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.
Link Errors (Sync Loss)	The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.
Port Congestion Index*	The estimated degree to which frame transmission was delayed due to a lack of buffer credits. This value is generally 0 - 100. The value 0 means there was no congestion. The value can exceed 100 if the buffer credit exhaustion persisted for an extended amount of time. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.
Port Protocol Errors (Port Send Delay Time)	The average number of milliseconds of delay that occur on the port for each send operation. The reason for these delays might be a lack of buffer credits. You cannot view zero buffer credit performance metrics for 16 Gbps Fibre Channel ports on resources that run IBM Spectrum Virtualize. Use the Port Send Delay Time metric if the Zero Buffer Credit Timer metric is not available.
Port Protocol Errors (Port Send Delay I/O Percentage)	The percentage of send operations where a delay occurred, relative to the total number of send operations that were measured for the port. Use this metric with the Port Send Delay Time metric to distinguish a few long delays from many short delays.

Table 151. Error rate metrics for ports (continued)	
Metric	Definition
Port Protocol Errors (Zero Buffer Credit Percentage)	The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that the node was reset. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.
Port Protocol Errors (Zero Buffer Credit Timer)	The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that the node was reset. If this metric is not available, use the Port Send Delay Time metric instead.

Note: *The performance metric for Port Congestion Index is only available for ports.

Table 152. Miscellaneous metrics for ports	
Metric	Definition
Port-to-Local Node Queue Time (Receive)	The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from other nodes that are in the local cluster.
Port-to-Local Node Queue Time (Send)	The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local cluster.
Port-to-Local Node Queue Time (Overall)	The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and other nodes that are in the local cluster.
Port-to-Remote Node Queue Time (Receive)	The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from a node that is in the remote cluster.
Port-to-Remote Node Queue Time (Send)	The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to a node that is in the remote cluster.

Table 152. Miscellaneous metrics for ports (continued)	
Metric	Definition
Port-to-Remote Node Queue Time (Overall)	The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and a node that is in the remote cluster.

Node performance metrics

Unless otherwise noted, you can view node metrics for the following resources:

- Nodes
- I/O Groups
- Storage systems

Table 153. Metrics for nodes	
Metric	Definition
Compression CPU Utilization (Core 1 to Core 28)	The approximate percentage of time that a processor core was busy with data compression tasks. The performance of each core is shown with a separate metric. Note that the value for this metric will be zero or close to zero if compression accelerator hardware is installed in the nodes.
CPU Utilization (Compression CPU)	The average percentage of time that the processors used for data compression I/O tasks are busy.
CPU Utilization (System CPU)	The average percentage of time that the processors on nodes are busy doing system I/O tasks.
Data Movement Rate (MiBs) ^{"1" on page 311}	The capacity, in MiBs per second, of the valid data in a reclaimed volume extent that garbage collection has moved to a new extent in the data reduction pool on the node. The valid data must be moved so that the whole extent can be freed up or reused to write new data. This metric corresponds to the collected mm statistic.
Data Rewrite Rate (MiBs) ^{"1" on page 311}	The rate, in MiBs per second, at which data is rewritten when a host overwrites data in data reduction pools on the node. The new version of the host data is written to a different location so that the capacity that was used by the previous version of the host data can be freed up and reclaimed. This metric corresponds to the collected cm statistic. You can view this metric for nodes only.

Table 153. Metrics for nodes (continued)	
Metric	Definition
Extent Collection Rate (cnt/s) ^{"1" on page 311}	The number of volume extents per second that were processed for garbage collection. The reclaimable capacity in the volume extents is collected so that it can be reused in the data reduction pools on the node. This metric corresponds to the collected ext col statistic. You can view this metric for nodes only.
Max Read Cache Fullness (%) ^{"2" on page 311}	The maximum amount of the lower cache which the cache partitions of the pools that are managed by the node are using for read operations. If the maximum value for the cache is 100%, the read cache partition for one or more of the pools is full. The read operations that pass through the node to the affected pools will be queued and the I/O response times will increase for the volumes in the affected pools. This metric corresponds to the collected lower cache rfmx statistic.
Max Write Cache Fullness (%) <u>"2" on page 311</u>	The maximum amount of the lower cache which the cache partitions of the pools that are managed by the node are using for write operations. If the maximum value for the cache is 100%, the write cache partition for one or more of the pools is full. The write operations that pass through the node to the affected pools will be queued and the I/O response times will increase for the volumes in the affected pools. This metric corresponds to the collected lower cache wfmx statistic.
New Address Write Rate (MiBs) <u>"1" on page 311</u>	The capacity in MiBs per second that was used to write the host's data to unallocated addresses in the data reduction pool on the node. Review this metric to determine which hosts are increasing the amount of capacity that is being written to data reduction pools on a node. This metric corresponds to the collected nm statistic. You can view this metric for nodes only.
Node Utilization by Node	The average of the bandwidth percentages of those ports in the node that are actively used for host and MDisk send and receive operations. The average is weighted by port speed and adjusted according to the technology limitations of the node hardware.
Read Cache Fullness (%) ^{"2" on page 311}	The average amount of the lower cache which the cache partitions of the pools that are managed by the node are using for read operations. Monitor the average cache fullness for read operations to identify the nodes that are experiencing heavy cache usage. This metric corresponds to the collected lower cache rfav statistic.

Table 153. Metrics for nodes (continued)	
Metric	Definition
Reclaimable Capacity (MiB) ^{"1" on page 311}	The capacity that can be reclaimed in the data reduction pools on the node. This metric corresponds to the collected rec statistic. You can view this metric for nodes only.
Recovered Capacity Rate (MiBs) ^{"1" on page 311}	The capacity in number of MiBs per second that was recovered by garbage collection for reuse in the data reduction pools on the node. This metric corresponds to the collected rm statistic. You can view this metric for nodes only.
System CPU Utilization (Core 1 to Core 28)	The approximate percentage of time that a processor core was busy with system I/O tasks. The performance of each core is shown with a separate metric.
Write Cache Fullness (%) <u>"2" on page 311</u>	The average amount of the lower cache which the cache partitions of the pools that are managed by the node are using for write operations. Monitor the average cache fullness for write operations to identify the nodes that are experiencing heavy cache usage. This metric corresponds to the collected lower cache wfav statistic.

- 1. This garbage collection metric applies to systems that are running IBM Spectrum Virtualize V8.1.2 or later. To view details about collected statistics, see Starting statistics collection.
- 2. This cache fullness metric applies to systems that are running IBM Spectrum Virtualize V7.3 or later.

 To view details about collected statistics, see Starting statistics collection.

Performance metrics for XIV, IBM Spectrum Accelerate, IBM FlashSystem A9000, and IBM FlashSystem A9000R

Monitor the performance metrics that are collected for XIV systems, IBM Spectrum Accelerate, IBM FlashSystem A9000, and IBM FlashSystem A9000R.

For XIV systems, IBM FlashSystem A9000, and IBM FlashSystem A9000R, performance metrics are available for the following resources:

- "Volume metrics" on page 311
- "Port metrics" on page 315

For IBM Spectrum Accelerate, performance metrics are available for the following resource:

• "Volume metrics" on page 311

Restriction: XIV systems, IBM Spectrum Accelerate, IBM FlashSystem A9000, and IBM FlashSystem A9000R do not track performance statistics for volumes that were never used. Because there are no performance statistics, performance metrics are not shown for these volumes and their related components.

Volume metrics

Volume performance metrics are organized into the following categories:

- Key storage metrics
- · Cache hit percentage metrics

- Response time metrics
- Miscellaneous metrics

Unless otherwise noted, you can view volume metrics for the following resources:

- Volumes
- Pools
- Modules
- · Host connections
- Storage systems

Table 154. Key volume metrics	
Metric	Description
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

 Metric
 Description

 Data Cache Hits (Read) "2" on page 313 the cache.
 The percentage of all read data that is read from the cache.

 Data Cache Hits (Write) "2" on page 313 the cache.
 The percentage of all write data that is written to

cache slots that are marked as modified.

Table 155. Cache hit percentage metrics (continued)	
Metric	Description
Data Cache Hits (Overall) ^{"2" on page 313}	The percentage of all data that is handled in the cache. This value includes read data that is read from the cache and write data that is written to cache slots that are marked as modified.
Overall I/O Cache Hits (Read) 11 on page 313	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.
Overall I/O Cache Hits (Write) <u>"1" on page 313</u>	The average percentage of all write operations that are handled in the cache, across all volumes on the server. This value includes both sequential and nonsequential write operations.
Overall I/O Cache Hits (Total) 10 n page 313	The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
SSD Read Cache Hits (I/O) ^{"2" on page 313}	The percentage of read operations that find data in the cache on a solid-state drive (SSD). The value for this metric is also included in the value for the Overall Read Cache Hit Percentage metric. This metric is only available for XIV systems that use solid-state drives as drives for caching.
SSD Read Cache Hits (Data) ^{"2" on page 313}	The percentage of all read data that was read from cache memory on a solid-state drive. This metric is only available for XIV systems that use solid-state drives as drives for caching.

- 1. For XIV systems, this metric is available in version 10.2.2 or later.
- 2. For XIV systems, this metric is available in version 10.2.4 or later.

Table 156. Response time metrics	
Metric	Description
Cache Hit Response Time (Read)	The average number of milliseconds to complete a read-cache hit operation.
Cache Hit Response Time (Write)	The average number of milliseconds to complete a write-cache hit operation.

Table 156. Response time metrics (continued)	
Metric	Description
Cache Hit Response Time (Overall)	The average number of milliseconds to complete a cache hit operation. This value includes the times for both read-cache hit and write-cache hit operations.
Cache Miss Response Time (Read)	The average number of milliseconds to complete a read-cache miss operation.
Cache Miss Response Time (Write)	The average number of milliseconds to complete a write-cache miss operation.
Cache Miss Response Time (Overall)	The average number of milliseconds to complete a cache miss operation. This value includes the times for both read-cache miss and write-cache miss operations.
Response Time by Transfer Size (Small)	The average number of milliseconds to complete an I/O operation that has a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.
Response Time by Transfer Size (Medium)	The average number of milliseconds to complete an I/O operation that has a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.
Response Time by Transfer Size (Large)	The average number of milliseconds to complete an I/O operation that has a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.
Response Time by Transfer Size (Very Large)	The average number of milliseconds to complete an I/O operation with a data transfer size that is greater than 512 KiB.
SSD Read Cache Hit Response Time	The average number of milliseconds that it takes to complete a hit operation on the read cache on a solid-state drive. The value for this metric is also included in the value for the Read Cache Hit Response Time metric. This metric is only available for XIV systems that use solid-state drives as drives for caching.

Note: For XIV systems, all the response time metrics are available for version 10.2.2 or later.

Table 157. Miscellaneous metrics	
Metric	Description
Average Transfer Size (Read) ^{"1" on page 315}	The average number of KiB that are transferred per read operation.
Average Transfer Size (Write) "1" on page 315	The average number of KiB that are transferred per write operation.
Average Transfer Size (Overall) <u>"1" on page 315</u>	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Table 157. Miscellaneous metrics (continued)	
Metric	Description
I/O Transfer Size (Small) "2" on page 315	The percentage of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.
I/O Transfer Size (Medium) ^{"2" on page 315}	The percentage of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.
I/O Transfer Size (Large) ^{"2" on page 315}	The percentage of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.
I/O Transfer Size (Very Large) "2" on page 315	The percentage of I/O operations with a data transfer size that is greater than 512 KiB.
Data Transfer Size (Small) "2" on page 315	The percentage of data that is transferred as a result of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.
Data Transfer Size (Medium) "2" on page 315	The percentage of data that is transferred as a result of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.
Data Transfer Size (Large) ^{"2" on page 315}	The percentage of data that is transferred as a result of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.
Data Transfer Size (Very Large) "2" on page 315	The percentage of data that is transferred as a result of I/O operations with a data transfer size that is greater than 512 KiB.
Volume Utilization ^{"1" on page 315} , "3" on page 315	The average percentage of time that the volume is busy.
Pool Activity Score "1" on page 315, "3" on page 315	The activity level of pools, which is set to the following value: [Read I/O Rate × (1- Read I/O Cache Hit %)] ÷ Total Pool Capacity
	The value of this metric is the same for all pools in the storage system.

- 1. For XIV systems, this metric is available in version 10.2.2 or later.
- 2. For XIV systems, this metric is available in version 10.2.4 or later.
- 3. This metric is only available when you view the performance of pools.

Port metrics

The following performance metrics for ports are available for XIV systems that are version 10.2.4 or later:

• I/O rate metrics

- Data rate metrics
- Response time metrics
- Miscellaneous port metrics

Restriction: Port metrics are not available for IBM Spectrum Accelerate.

Table 158. I/O rate metrics	
Metric	Description
I/O Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.
I/O Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.
I/O Rate (Total)	The average number of send operations and receive operations per second.

Table 159. Data rate metrics	
Metric	Description
Data Rate (Send)	The average rate at which data is sent through the port. The rate is measured in MiB per second.
Data Rate (Receive)	The average rate at which data is received by the port. The rate is measured in MiB per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Table 160. Response time metrics	
Metric	Description
Response Time (Send)	The average number of milliseconds to complete a send operation.
Response Time (Receive)	The average number of milliseconds to complete a receive operation.
Response Time (Overall)	The average number of milliseconds to complete a send or receive operation.

Table 161. Miscellaneous port metrics	
Metric	Description
Bandwidth (Send)	The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth (Receive)	The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Table 161. Miscellaneous port metrics (continued)	
Metric	Description
Bandwidth (Overall)	The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Performance metrics for IBM Spectrum Scale

Monitor the performance metrics that are collected for IBM Spectrum Scale storage systems.

Overview

The performance metrics are divided into the following categories:

- "Node performance metrics" on page 317
- "File system performance metrics" on page 317

Node performance metrics

You can view the following metrics for each IBM Spectrum Scale cluster node.

Table 162. Metrics for nodes	
Metric	Definition
CPU Utilization (User)	The average percentage of time that the processors on a node are busy doing user tasks.
CPU Utilization (System)	The average percentage of time that the processors on a node are busy doing system tasks.
CPU Utilization (Total)	The average percentage of time that the processors on a node are busy doing user tasks and system tasks.
Memory Used (Cache and Buffer)	The average percentage of memory on a node that is used for cache and buffer memory.
Memory Used	The average percentage of memory that is used on a node. This value does not include the memory that is used for cache and buffer memory.
Memory Used (Total)	The average percentage of memory that is used on a node. This value includes the memory that is used for cache and buffer memory.
I/O Rate (Read)	The average number of read operations per second.
I/O Rate (Write)	The average number of write operations per second.
I/O Rate (Total)	The average number of read operations and write operations per second.

File system performance metrics

You can view file system metrics for the following resources:

- File systems
- IBM Spectrum Scale storage systems

Tip: For a storage system, the metrics contain summary values for all the file systems on the storage system.

Table 163. Metrics for file systems	
Metric	Definition
Data Rate (Read)	The average number of MiB per second that are transferred for read operations.
Data Rate (Write)	The average number of MiB per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Maximum Data Rate (Read)	The maximum number of MiB per second that are transferred for read operations.
Maximum Data Rate (Write)	The maximum number of MiB per second that are transferred for write operations.
Maximum Data Rate (Total)	The maximum number of MiB per second that are transferred for read operations and write operations.
I/O Rate (Read)	The average number of read operations per second.
I/O Rate (Write)	The average number of write operations per second.
I/O Rate (Total)	The average number of read operations and write operations per second.
Maximum I/O Rate (Read)	The maximum number of read operations per second.
Maximum I/O Rate (Write)	The maximum number of write operations per second.
Maximum I/O Rate (Total)	The maximum number of read operations and write operations per second.
Response Time (Read)	The average number of milliseconds for the backend storage resources to respond to a read operation from the file system.
Response Time (Write)	The average number of milliseconds for the backend storage resources to respond to a write operation from the file system.
Response Time (Overall)	The average number of milliseconds for the backend storage resources to respond to a read operation or a write operation from the file system.

Performance metrics for IBM FlashSystem 840 and IBM FlashSystem 900

To review trends in performance for IBM FlashSystem 840 and IBM FlashSystem 900 storage systems, you add performance metrics to performance charts. Use the charts to monitor the performance of the storage systems.

Overview

The performance metrics are described in the following sections:

- "Volume metrics" on page 319
- "Disk metrics" on page 320
- "Port metrics" on page 320

Restrictions:

- IBM FlashSystem 840 and IBM FlashSystem 900 storage systems do not track performance for individual volumes. The volume metrics that are available are provided to allow comparisons across different storage systems, but represent performance that is measured for ports rather than for volumes.
- IBM Spectrum Control supports only Fibre Channel adapters and ports that are installed on storage systems. For IBM FlashSystem 840 and IBM FlashSystem 900 storage systems that have other types of ports, performance data is only available for the drives.

Volume metrics

IBM FlashSystem 840 and IBM FlashSystem 900 storage systems do not track the performance of individual volumes. Therefore, the performance metrics that you can view represent the sum of the measured performance of all the ports in the storage system.

Table 164. Metrics for volumes	
Metric	Description
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Total)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Disk metrics

You can view the following performance metrics for drives or flash modules on IBM FlashSystem 840 and IBM FlashSystem 900 storage systems.

Table 165. Metrics for disks	
Metric	Description
Data Rate (Read)	The average number of MiB per second that are read from the back-end storage resources.
Data Rate (Write)	The average number of MiB per second that are written to the back-end storage resources.
Data Rate (Total)	The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.
Flash Health Percentage	The overall flash module health. The health percentage is calculated based on the number of unusable blocks on the flash module.

Port metrics

You can view the following port performance metrics for Fibre Channel ports on IBM FlashSystem 840 and IBM FlashSystem 900 storage systems. No performance data is available for InfiniBand or FCoE ports.

Table 166. Metrics for ports	
Metric	Description
Bandwidth (Send)	The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth (Receive)	The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth (Total)	The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Data Rate (Send)	The average rate at which data is sent through the port. The rate is measured in MiB per second.
Data Rate (Receive)	The average rate at which data is received by the port. The rate is measured in MiB per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
I/O Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.

Table 166. Metrics for ports (continued)	
Metric	Description
I/O Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.
I/O Rate (Total)	The average number of send operations and receive operations per second.
Response Time (Send)	The average number of milliseconds to complete a send operation.
Response Time (Receive)	The average number of milliseconds to complete a receive operation.
Response Time (Overall)	The average number of milliseconds to complete a send or receive operation.
Transfer Size (Send)	The average number of KiB that are transferred per send operation.
Transfer Size (Receive)	The average number of KiB that are transferred per receive operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Performance metrics for Dell EMC storage systems

Monitor the performance metrics that are collected for Unity, VMAX, and VNX storage systems.

Overview

The performance metrics are described in the following sections:

- "Unity storage systems" on page 321
- "VMAX storage systems" on page 326
- "VNX storage systems" on page 330

Restriction: For Dell EMC, performance metrics are only available for block storage systems.

Unity storage systems

The following performance metrics are available for Unity resources:

- "Volume metrics for Unity" on page 321
- "Disk metrics for Unity" on page 323
- "Port metrics for Unity" on page 324
- "Node metrics for Unity" on page 325
- "File system metrics for Unity" on page 325

A performance metric might apply to one or more storage resources. To check which resources the performance metric applies to, see the tips and the table footnotes.

Volume metrics for Unity

Volume performance metrics are divided into the following categories:

- Key volume metrics
- · Cache hit metrics

Tip: You can view the volume metrics in Key volume metrics for the following resources:

- Volumes
- Pools
- Nodes
- Storage systems

Table 167. Key volume metrics	
Description	
The average number of read operations per second. This value includes both sequential and nonsequential read operations.	
The average number of write operations per second. This value includes both sequential and nonsequential write operations.	
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.	
The average number of MiBs per second that are transferred for read operations.	
The average number of MiBs per second that are transferred for write operations.	
The average number of MiB per second that are transferred for read operations and write operations.	
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.	
The average number of KiB that are transferred per read operation.	
The average number of KiB that are transferred per write operation.	
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.	

Tip: You can view the volume metrics in Cache hit metrics for the following resources:

- Nodes
- Storage systems

Table 168. Cache hit metrics	
Metric	Description
Overall I/O Cache Hits (Read) %	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.
Overall I/O Cache Hits (Write) %	The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.
Overall I/O Cache Hits (Total) %	The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Disk metrics for Unity

You can view the metrics in <u>Key disk metrics</u> for the following resources:

- Pools
- Disks
- Nodes
- Storage systems

Table 169. Key disk metrics	
Metric	Definition
I/O Rate (Read)	The average number of read operations per second that are issued to the back-end storage resources.
I/O Rate (Write)	The average number of write operations per second that are issued to the back-end storage resources.
I/O Rate (Total)	The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.
Data Rate (Read)	The average number of MiB per second that are read from the back-end storage resources.
Data Rate (Write)	The average number of MiB per second that are written to the back-end storage resources.

Table 169. Key disk metrics (continued)	
Metric	Definition
Data Rate (Total)	The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.
Response Time (Overall)	The average number of milliseconds for the backend storage resources to respond to a read or a write operation.
Transfer Size (Read)	The average number of KiB that are transferred per read operation from the back-end storage resources.
Transfer Size (Write)	The average number of KiB that are transferred per write operation to the back-end storage resources.
Transfer Size (Overall)	The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Port metrics for Unity

You can view the metrics in $\underline{\text{Key port metrics}}$ for ports, nodes, and storage systems.

Table 170. Key port metrics	
Metric	Description
I/O Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.
I/O Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.
I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Send)	The average rate at which data is sent through the port. The rate is measured in MiB per second.
Data Rate (Receive)	The average rate at which data is received by the port. The rate is measured in MiB per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
Transfer Size (Send)	The average number of KiB that are transferred per send operation.
Transfer Size (Receive)	The average number of KiB that are transferred per receive operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Node metrics for Unity

You can view the metrics in Key node metrics for nodes and storage systems.

Table 171. Key node metrics	
Metric	Description
System CPU	The average percentage of time that the processors on nodes are busy doing system I/O tasks.

File system metrics for Unity

Unless otherwise noted, you can view the performance metrics in Key file system metrics for the following resources:

- Pools
- Nodes
- · File systems
- Storage systems

Tip: For a storage system, the metrics contain summary values for all of the file systems on the storage system.

Table 172. Key file systems metrics	
Metric	Definition
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

VMAX storage systems

The following performance metrics are available for VMAX resources:

- "Volume metrics for VMAX" on page 326
- "Disk metrics for VMAX" on page 328
- "Port metrics for VMAX" on page 329

A performance metric might apply to one or more storage resources. To check which resources the performance metric applies to, see the tips and the table footnotes.

Volume metrics for VMAX

Volume performance metrics are divided into the following categories:

- Key volume metrics
- · Cache hit metrics
- Transfer size metrics

Tip: Unless otherwise noted, you can view the volume metrics in <u>Key volume metrics</u>, <u>Cache hit metrics</u>, and Transfer size metrics for the following resources:

- Volumes
- Directors
- · Storage systems

Table 173. Key volume metrics	
Metric	Description
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Response Time (Read) "1" on page 327	The average number of milliseconds to complete a read operation.
Response Time (Write) "1" on page 327	The average number of milliseconds to complete a write operation.
Response Time (Overall) 10 on page 327	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Table 173. Key volume metrics (continued)	
Metric	Description
Volume Utilization <u>"2" on page 327</u>	The average percentage of time that the volume is busy.
Normal I/O Rate (Read)	The average number of nonsequential read operations per second.
Normal I/O Rate (Write)	The average number of nonsequential write operations per second.
Normal I/O Rate (Total)	The average number of nonsequential I/O operations per second. This value includes both read and write operations.
Sequential I/O Rate (Read)	The average number of sequential read operations per second.
Sequential I/O Rate (Write)	The average number of sequential write operations per second.
Sequential I/O Rate (Total)	The average number of sequential I/O operations per second. This value includes both read and write operations.

Note:

- 1. This metric is only available when you view the performance of volumes or storage systems.
- 2. This metric is only available when you view the performance of volumes.

Table 174. Cache hit metrics	
Metric	Description
Normal I/O Cache Hits (Read)	The percentage of nonsequential read operations that find data in the cache. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Normal I/O Cache Hits (Write)	The percentage of nonsequential write operations that are handled in the cache.
Normal I/O Cache Hits (Total)	The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Sequential I/O Cache Hits (Read)	The percentage of sequential read operations that find data in the cache. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Table 174. Cache hit metrics (continued)	
Metric	Description
Sequential I/O Cache Hits (Write)	The percentage of sequential write operations that are handled in the cache.
Sequential I/O Cache Hits (Total)	The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Overall I/O Cache Hits (Read)	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.
Overall I/O Cache Hits (Write)	The average percentage of all write operations that are handled in the cache, across all volumes on the server. This value includes both sequential and nonsequential write operations.
Overall I/O Cache Hits (Total)	The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Table 175. Transfer size metrics	
Metric	Description
Average Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Average Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Average Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Disk metrics for VMAX

Disk performance metrics for VMAX block storage systems are divided into the following categories:

- Key disk metrics
- Transfer size metrics

Table 176. Key disk metrics	
Metric	Definition
I/O Rate (Read)	The average number of read operations per second that are issued to the back-end storage resources.
I/O Rate (Write)	The average number of write operations per second that are issued to the back-end storage resources.
I/O Rate (Total)	The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.
Data Rate (Read)	The average number of MiB per second that are read from the back-end storage resources.
Data Rate (Write)	The average number of MiB per second that are written to the back-end storage resources.
Data Rate (Total)	The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.
Response Time (Overall)	The average number of milliseconds for the backend storage resources to respond to a read or a write operation.
Disk Utilization Percentage	The approximate utilization percentage of a particular array over a particular time interval, that is, the average percent of time that the disk associated with the array was busy.

Table 177. Transfer size metrics	
Metric	Description
Transfer Size (Read)	The average number of KiB that are transferred per read operation from the back-end storage resources.
Transfer Size (Write)	The average number of KiB that are transferred per write operation to the back-end storage resources.
Transfer Size (Overall)	The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Port metrics for VMAX

You can view the metrics below for ports and storage systems.

Table 178. Port metrics	
Metric	Description
I/O Rate (Total)	The average number of send operations and receive operations per second.

Table 178. Port metrics (continued)	
Metric	Description
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
Response Time (Overall)	The average number of milliseconds to complete a send or receive operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

VNX storage systems

The following performance metrics are available for VNX resources:

- "Volume metrics for VNX" on page 330
- "Disk metrics for VNX" on page 331
- "Port metrics for VNX" on page 332

Volume metrics for VNX

Volume performance metrics are divided into the following categories:

- Key volume metrics
- Transfer size metrics

Tip: Unless otherwise noted, you can view the volume metrics in <u>Key volume metrics</u> and <u>Transfer size</u> metrics for the following resources:

- Volumes
- Nodes
- Storage systems

Table 179. Key volume metrics	
Metric	Description
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.

Table 179. Key volume metrics (continued)	
Metric	Description
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.

Table 180. Transfer size metrics	
Metric	Description
Average Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Average Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Average Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Disk metrics for VNX

Disk performance metrics for VNX storage systems are divided into the following categories:

- Key disk metrics
- Transfer size metrics

Table 181. Key disk metrics	
Metric	Definition
I/O Rate (Read)	The average number of read operations per second that are issued to the back-end storage resources.
I/O Rate (Write)	The average number of write operations per second that are issued to the back-end storage resources.
I/O Rate (Total)	The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.
Data Rate (Read)	The average number of MiB per second that are read from the back-end storage resources.
Data Rate (Write)	The average number of MiB per second that are written to the back-end storage resources.
Data Rate (Total)	The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Table 182. Transfer size metrics	
Metric	Description
Transfer Size (Read)	The average number of KiB that are transferred per read operation from the back-end storage resources.
Transfer Size (Write)	The average number of KiB that are transferred per write operation to the back-end storage resources.
Transfer Size (Overall)	The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Port metrics for VNX

You can view the metrics below for ports and storage systems.

Table 183. Port metrics	
Metric	Description
I/O Rate (Total)	The average number of send operations and receive operations per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Performance metrics for Hitachi VSP storage systems

Monitor the performance metrics that are collected for storage systems.

The following performance metrics are available for Hitachi VSP storage systems.

- "Volume metrics" on page 332
- "Port metrics" on page 334
- "Node metrics" on page 334

A performance metric might apply to one or more storage resources. To check which resources the performance metric applies to, see the tips.

Volume metrics

Volume performance metrics are divided into the following categories:

- Key volume metrics
- · Cache hit metrics

Tip: You can view the volume metrics in <u>Key volume metrics</u> for the following resources:

- Volumes
- Pools
- Storage systems

Table 184. Key volume metrics	
Metric	Description
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Tip: You can view the volume metrics in <u>Cache hit metrics</u> for the following resources:

- Volumes
- Pools
- Storage systems

Table 185. Cache hit metrics	
Metric	Description
Overall I/O Cache Hits (Read) %	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.
Overall I/O Cache Hits (Write) %	The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Port metrics

You can view the metrics in Key port metrics for ports and storage systems.

Table 186. Key port metrics	
Metric	Description
I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
Response Time (Overall)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Node metrics

You can view the metrics in Key node metrics for storage systems.

Table 187. Key node metrics	
Metric	Description
System CPU	The average percentage of time that the processors on nodes are busy doing system I/O tasks.

Performance metrics for NetApp storage systems

Monitor the performance metrics that are collected for NetApp storage systems that are running ONTAP o

The following performance metrics are available for storage systems that are running ONTAP 9:

- Volume metrics
- · Disk metrics
- Port metrics
- Node metrics
- File system metrics

A performance metric might apply to one or more storage resources. To check which resources the performance metric applies to, see the tips.

Volume metrics

Volume performance metrics are divided into the following categories:

- Key volume metrics
- Cache hit metrics

Tip: You can view the volume metrics in Key volume metrics for the following resources:

- Volumes
- Pools
- Nodes
- Storage systems

Table 188. Key volume metrics	
Metric	Description
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.

Table 188. Key volume metrics (continued)	
Metric	Description
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Volume cache hit metrics are only available when you view the performance of nodes or storage systems. To check which resources a volume cache hit metric applies to, see the table footnotes.

Table 189. Cache hit metrics	
Metric	Description
Overall I/O Cache Hits (Read) % ^{"1" on page 336}	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Notes:

1. This metric is only available when you view the performance of storage systems.

Disk metrics

Disk performance metrics for NetApp storage systems that are running ONTAP 9 are described in <u>Key disk</u> metrics.

Tip: You can view the disk metrics in Key disk metrics for the following resources:

- Pools
- Disks
- Nodes
- Storage systems

Table 190. Key disk metrics	
Metric	Definition
I/O Rate (Read)	The average number of read operations per second that are issued to the back-end storage resources.
I/O Rate (Write)	The average number of write operations per second that are issued to the back-end storage resources.

Table 190. Key disk metrics (continued)		
Metric	Definition	
I/O Rate (Total)	The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.	
Data Rate (Read)	The average number of MiB per second that are read from the back-end storage resources.	
Data Rate (Write)	The average number of MiB per second that are written to the back-end storage resources.	
Data Rate (Total)	The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.	
Response Time (Read)	The average number of milliseconds for the backend storage resources to respond to a read operation.	
Response Time (Write)	The average number of milliseconds for the backend storage resources to respond to a write operation.	
Response Time (Overall)	The average number of milliseconds for the backend storage resources to respond to a read or a write operation.	
Transfer Size (Read)	The average number of KiB that are transferred per read operation from the back-end storage resources.	
Transfer Size (Write)	The average number of KiB that are transferred per write operation to the back-end storage resources.	
Transfer Size (Overall)	The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.	

Port metrics

Port performance metrics are described in $\underline{\text{Key port metrics}}.$

Tip: You can view the port metrics in Key port metrics for the following resources:

- Ports
- Nodes
- Storage systems

Table 191. Key port metrics	
Metric	Description
I/O Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.
I/O Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.

Table 191. Key port metrics (continued)	
Metric	Description
I/O Rate (Total)	The average number of send operations and receive operations per second.
Response Time (Send)	The average number of milliseconds to complete a send operation.
Response Time (Receive)	The average number of milliseconds to complete a receive operation.
Response Time (Overall)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
Data Rate (Send)	The average rate at which data is sent through the port. The rate is measured in MiB per second.
Data Rate (Receive)	The average rate at which data is received by the port. The rate is measured in MiB per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
Transfer Size (Send)	The average number of KiB that are transferred per send operation.
Transfer Size (Receive)	The average number of KiB that are transferred per receive operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Node metrics

You can view the metrics in Key node metrics for storage systems.

Table 192. Key node metrics	
Metric	Description
System CPU	The average percentage of time that the processors on nodes are busy doing system I/O tasks.

File system metrics

Unless otherwise noted, you can view the performance metrics in Key file system metrics for the following resources:

- Pools
- Nodes
- File systems
- Storage systems

Tip: For a storage system, the metrics contain summary values for all of the file systems on the storage system.

Table 193. Key file systems metrics	
Metric	Definition
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Performance metrics for Pure storage systems

Monitor the performance metrics that are collected for your Pure FlashArray//M and FlashArray//X storage systems.

A performance metric might apply to one or more storage resources. To check which resources the performance metric applies to, see the tips.

Volume metrics

Volume performance metrics are described in the Key volume metrics table.

Tip: You can view the volume metrics in Key volume metrics for the following resources:

- Volumes
- Storage systems

Table 194. Key volume metrics	
Metric	Description
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Performance metrics for other storage systems

Monitor the performance metrics that are collected for SMI-S BSP systems such as DS5000, DS4000, and DS3000, , and other storage systems.

Overview

Performance metrics are available for the following resources:

- "Volume metrics" on page 340
- "Port metrics" on page 342

A performance metric might apply to one or more storage resources. To check which resources the performance metric applies to, see the tips and the table footnotes.

Restriction: For NetApp storage systems managed by SMI-S providers, performance monitoring is supported only for volumes and front-end ports. It is not supported at the storage system level.

Volume metrics

Volume performance metrics are divided into the following categories:

- Key volume metrics
- · Cache hit percentage metrics
- Average transfer size metrics

Tip: Unless otherwise noted, you can view the volume metrics in <u>Key volume metrics</u>, <u>Cache hit metrics</u>, and <u>Average transfer size metrics</u> for the following resources:

- Volumes
- Nodes
- Storage systems

Table 195. Key volume metrics	
Metric	Description
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
Volume Utilization <u>"1" on page 341</u>	The average percentage of time that the volume is busy.

Note:

1. This metric is only available when you view the performance of volumes.

Table 196. Cache hit percentages	
Metric	Description
Overall I/O Cache Hits (Read)	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Table 196. Cache hit percentages (continued)	
Metric	Description
Overall I/O Cache Hits (Write)	The average percentage of all write operations that are handled in the cache, across all volumes on the server. This value includes both sequential and nonsequential write operations.
Overall I/O Cache Hits (Total)	The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Table 197. Average transfer size metrics	
Metric	Description
Average Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Average Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Average Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Port metrics

Port performance metrics are divided into the following categories:

- Key port metrics
- Transfer size metrics

Tip: You can view the port metrics in $\underline{\text{Key port metrics}}$ and $\underline{\text{Transfer size metrics}}$ for ports and storage systems.

Table 198. Key port metrics	
Metric	Description
I/O Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.
I/O Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.
I/O Rate (Total)	The average number of send operations and receive operations per second.
Data Rate (Receive)	The average rate at which data is received by the port. The rate is measured in MiB per second.
Data Rate (Send)	The average rate at which data is sent through the port. The rate is measured in MiB per second.

Table 198. Key port metrics (continued)	
Metric	Description
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Table 199. Transfer size metrics	
Metric	Description
Average Transfer Size (Receive)	The average number of KiB that are transferred per receive operation.
Average Transfer Size (Send)	The average number of KiB that are transferred per send operation.
Average Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Performance metrics for switches

Monitor the performance metrics that are collected for physical switches, switch ports, and inter-switch connections.

Performance metrics for switches are divided into the following categories:

- Key port metrics
- I/O rate metrics
- Peak data rate metrics
- Frame error rate metrics
- Port protocol error rate metrics
- Link error rate metrics
- Miscellaneous metrics

Table 200. Key port metrics	
Metric	Description
Bandwidth Percentage (Send)	The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth Percentage (Receive)	The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth Percentage (Overall)	The percentage of the port bandwidth that is used for send operations and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Data Rate (Send)	The average rate at which data is sent by the port. A send operation is a read operation that is processed, or a write operation that is initiated by the port. The rate is measured in MiB per second.

Table 200. Key port metrics (continued)	
Metric	Description
Data Rate (Receive)	The average rate at which data is received by the port. A receive operation is a write operation that is processed, or a read operation that is initiated by the port. The rate is measured in MiB per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
Total Port Error Rate	The average number of times per second that an error was detected on the port. This rate is a summation of all the other error rates for the port.

Table 201. I/O rates	
Metric	Description
Port Frame Rate (Send)	The average number of frames per second that are sent through the port.
Port Frame Rate (Receive)	The average number of frames per second that are received by the port.
Port Frame Rate (Total)	The average number of frames per second that are transferred. This value includes frames that are sent and frames that are received by the port.

Table 202. Peak data rates	
Metric	Description
Peak Data Rate (Send) ^{"1" on page 344}	The highest rate at which data is sent by the port. A send operation is a read operation that is processed, or a write operation that is initiated by the port.
Peak Data Rate (Receive) ^{"1" on page 344}	The highest rate at which data is received by the port. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Notes:

1. This metric is only available for ports on Brocade switches. You must use an SMI agent to collect data for this metric. In the Storage Management Initiative Specification (SMI-S), peak data rates for a port are reported in the CIM_FcPortRateStatistics class. For SMI-S Version 1.1, this is an optional class, meaning that it might not be supported by a given vendor. This can result in IBM Spectrum Control displaying zeros for the peak rates, even when there is non-zero traffic flowing through the port. When a vendor does support this class, the peak data rate represents the peak value of the associated metric counter (as reported by the CIM agent used to access the device) over a recent, vendor-defined measurement window (for example, 20 milliseconds) of the IBM Spectrum Control measurement window (for example, 15 minutes).

Table 203. Frame error rates	
Metric	Description
Bad EOF CRC Error Rate ^{"1" on page 345}	The percentage of nonsequential read operations that find data in the cache. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
CRC Error Rate	The percentage of nonsequential write operations that are handled in the cache.
Discarded Class 3 Frame Rate	The average number of class 3 frames per second that are discarded.
Error Frame Rate "1" on page 345	The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.
F-BSY Frame Rate ^{"2" on page 345}	The average number of F-BSY frames per second that are generated. An F-BSY frame is issued by the fabric to indicate that a frame cannot be delivered because the fabric or destination N_port is busy.
F-RJT Frame Rate ^{"2" on page 345}	The average number of F-RJT frames per second that are generated.An F-RJT frame is issued by the fabric to indicate that delivery of a frame was denied.
Long Frame Rate	The average number of frames that are received per second that are longer than 2140 octets. This number excludes start-of-frame bytes and end-of-frame bytes. The 2140 octet limit is calculated based on the assumption that a frame has 24 header bytes, 4 CRC bytes, and 2112 data bytes.
Short Frame Rate ^{"2" on page 345}	The average number of frames that are received per second that are shorter than 28 octets. This number excludes start-of-frame bytes and end-of-frame bytes. The 28 octet limit is calculated based on the assumption that a frame has 24 header bytes, and 4 CRC bytes.

Notes:

- 1. This metric is only available for ports on Brocade switches.
- 2. This metric is only available for ports on Cisco switches.

Table 204. Port protocol error rates	
Metric	Description
Class 3 Receive Timeout Frame Rate ^{"1" on page 347}	The average number of class 3 frames per second that were discarded after reception because of a timeout condition. The timeout condition occurs while a transmitting port waits for buffer credit from a port at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.
Class 3 Send Timeout Frame Rate ^{"1" on page 347}	The average number of class 3 frames per second that were discarded before transmission because of a timeout condition. The timeout condition occurs while the switch or port waits for buffer credit from the receiving port at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.
Credit Recovery Link Reset Rate	The estimated average number of link resets per second that a switch or port completed to recover buffer credits. This estimate attempts to disregard link resets that were caused by link initialization. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.
Discarded Frame Rate <u>"1" on page 347</u>	The average number of frames per second that are discarded because host buffers are unavailable for the port.
Link Reset Received Rate	The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state.
Link Reset Transmitted Rate	The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state.
Port Congestion Index	The estimated degree to which frame transmission was delayed due to a lack of buffer credits. This value is generally 0 - 100. The value 0 means there was no congestion. The value can exceed 100 if the buffer credit exhaustion persisted for an extended amount of time. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.
Zero Buffer Credit Percentage	The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that the node was reset. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.

Table 204. Port protocol error rates (continued)	
Metric	Description
Zero Buffer Credit Rate	The average number of Zero Buffer Credit conditions per second that occurred. A Zero Buffer Credit condition occurs when a port is unable to send frames because of a lack of buffer credit since the last node reset. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

Notes:

1. This metric is available only for ports on Brocade switches.

Table 205. Link error rates	
Metric	Description
Encoding Disparity	The average number of disparity errors per second that are received.
Invalid Link Transmission Rate	The average number of times per second that an invalid transmission word was detected by the port while the link did not experience any signal or synchronization loss.
Invalid Word Transmission Rate	The average number of bit errors per second that are detected.
Link Failure Rate	The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.
Loss of Signal Rate	The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.
Loss of Sync Rate	The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.
Primitive Sequence Protocol Error Rate	The average number of primitive sequence protocol errors per second that are detected. This error occurs when there is a link failure for a port.

Table 206. Miscellaneous port metrics	
Metric	Description
Link Quality Percentage	The percentage is based on whether the port is an expansion port (E_port) or a fabric port (F_port), and on the numbers and types of errors that are detected by the port.
Port Frame Size (Overall)	The average frame transfer size. This value is measured in KiB and includes frames that are sent and frames that are received by the port.
Port Frame Size (Receive)	The average size of a frame, in KiB, that is received by the port.
Port Frame Size (Send)	The average size of a frame, in KiB, that is sent through the port.

Monitoring the capacity of resources

IBM Spectrum Control can collect information about the capacity and space usage of block, file, and object storage resources. This information includes key metrics that can help you measure, identify, and troubleshoot capacity and space usage issues in your storage. You can view capacity metrics for storage resources, including tiers.

You can use the capacity and space usage information to complete the following tasks:

- Measure, compare, and troubleshoot the capacity and space usage of resources such as:
 - Block storage systems and pools
 - Compressed, thin-provisioned, and Easy Tier volumes
 - File systems, pools, and filesets in file storage systems
 - Containers in object storage systems
 - Tiers
- Review the alerts that were triggered when the capacity or space usage of a resource fell outside a specific range.
- Customize views of capacity and space usage so that you can analyze specific resources and metrics during time ranges that you specify.
- View capacity and space usage information in a chart or table format to help you quickly identify where and when issues are occurring. The chart is a visual representation of how the capacity and space usage of resources trend over time.
- Drill down into resources to view detailed information about the capacity and space usage of internal resources. For example, if a SAN Volume Controller is shown in the chart, you can quickly view and compare the capacity and space usage of its pools.
- Implement server-centric monitoring of SAN resources without requiring a Storage Resource agent. When you add an agentless server, IBM Spectrum Control correlates the server with the known host connections on monitored storage systems. If matches are found, you can view the amount of storage that is assigned to the server, and trace that storage back to the storage system. You can then view details about the internal resources of the related storage system, including capacity and space usage information.
- Export capacity and space usage information to a CSV file. A CSV file is a file that contains commadelimited values and can be viewed with a text editor or imported into a spreadsheet application.

Before you begin

Before you can view capacity and space usage information for resources, you must complete the following tasks:

Task	1 Learn more
Add storage systems for monitoring by IBM Spectrum Control.	For information about how to add these resources, see "Adding storage systems" on page 118.
Ensure that capacity data is being collected for a resource. Typically, data collection is scheduled when resources are added for monitoring.	For information about verifying data collection, see "Collecting data" on page 138.
Define capacity and capacity usage alerts to be notified if the capacity or capacity usage of a resource falls outside a specified range and might represent a potential problem.	For information about how to define capacity and capacity usage alerts, see "Defining alert definitions for general attributes and capacity changes" on page 159.
When you define a capacity or capacity usage alert for an alert policy or for an individual resource, select a specific metric that you want to measure. For example, you can define an alert that notifies you when the total used capacity for a storage system falls outside a specified range.	

Related reference

Capacity metrics for block storage systems

To review trends in capacity and space usage for storage, you add metrics to capacity charts. You use the charts to detect capacity shortages and space usage trends.

Capacity metrics for file storage systems

To review trends in capacity and space usage for file storage systems, you add metrics to capacity charts. You use the charts for filesets, file systems, and file system pools to detect capacity shortages and space usage trends.

Capacity metrics for object storage systems

To review trends in capacity and space usage for object storage systems, you add metrics to capacity charts. Use the charts for containers to detect capacity shortages and space usage trends for the containers in your object storage systems.

Capacity metrics for tiers

Review trends in the capacity and space usage for tiers. You use the charts to detect capacity shortages and space usage for the tiers in your storage environment.

Setting capacity limits

If your company has a policy to maintain a reserve of available capacity, you can set a limit on the amount of capacity that is used by your storage systems and pools.

Only users with Administrator privileges can set capacity limits.

When you set a limit on the amount of capacity that is used, you can monitor whether the capacity that is used is above or below the limit that you set. You can also define alerts that notify you when you are over the capacity limit.

You can set the same capacity limit:

- For all storage systems or one or more storage systems
- For all pools or one or more pools

Alternatively, you can set different capacity limits for storage systems and pools.

Tip: FlashSystem A9000 and FlashSystem A9000R storage systems are configured to report storage consumption at the storage system rather than the pool level. So, if you want to set a capacity limit for these types of storage systems, set the capacity limit for the storage systems.

1. To set the capacity limit for storage systems or pools, choose one of the following options:

Storage Resource	Step
Block Storage Systems	Click Storage , and then click Block Storage Systems .
Pools	Click Storage , and then click Pools .

2. Select the storage systems or pools.

Tip: To select multiple storage systems or pools, press Shift and click.

- 3. Right-click and click View Properties.
- 4. Click Edit Properties
- 5. Enter the percentage value for the capacity limit.
- 6. Save your changes.

The capacity limit is set and can be shown as a percentage value or GiB value. Just right-click any column heading and add Capacity Limit (%) or Capacity Limit (GiB) to the Block Storage Systems or Pools pages. You can also add the Adjusted Used Capacity (%) column which tells you how much used capacity is left based on the capacity limit. By default, the Capacity-to-Limit (GiB) column is shown on the Block Storage Systems and Pools pages so that you can see how much available capacity you have left before you reach the capacity limit.

Zero capacity: When you set the capacity limit for pools, the values shown for **Zero Capacity** are readjusted to take into account the capacity limit of the pool. The date will represent when the capacity limit of the pool is reached. If the pool has already reached the capacity limit, Depleted is shown. None is shown when a trend in storage consumption can't be detected because the pool's storage isn't being consumed or because not enough data was collected to predict storage consumption.

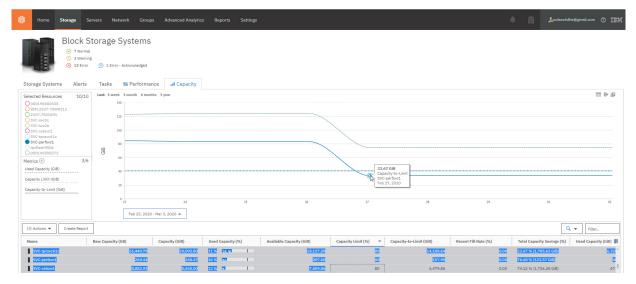
Scenario: Investigating compliance with the capacity limit

Track the compliance of your block storage systems with the capacity limit.

In this scenario, you set the capacity limit for the storage systems on your production platform to 80%. Today, you want to create a chart that plots the compliance of your block storage systems with the capacity limit that was set.

Because you want to track capacity in relation to the capacity limit, you remove the **Capacity (GiB)** metric from the chart. And, you add the following metrics to track whether the available capacity of each storage system is over or under the capacity limit:

- Capacity Limit (GiB)
- Capacity-to-Limit (GiB)
- 1. From the Storage menu, click Block Storage Systems.
- 2. Click the Capacity tab.
- Select the storage systems that you want to add to the capacity chart.
 To select multiple storage systems, press shift and click. You can add up to 10 storage systems at a time to the capacity chart.
- 4. Click the Select Chart Metrics button, add Capacity Limit (GiB) Capacity-to-Limit (GiB), and click Save.



Tip: To see the capacity metrics for a single resource, click a storage system in the **Selected Resources** section of the chart. And, if you want to focus on a specific period, you can change the default date range of a month to a week, or select start and end dates.

You can identify the storage systems that comply with the capacity limit that was set.

Removing the capacity limit

Remove the capacity limit that was set for your storage systems and pools.

Only users with Administrator privileges can remove the capacity limit.

1. To remove the capacity limit for storage systems or pools, choose one of the following options:

Storage Resource	Step
Block Storage Systems	Click Storage , and then click Block Storage Systems .
Pools	Click Storage , and then click Pools .

2. Select the storage systems or pools.

Tip: To select multiple storage systems or pools, press shift and click.

- 3. Right-click, and then click View Properties.
- 4. Click Edit.
- 5. Complete one of the following action:
 - If the pools or storage systems have the same capacity limit, delete the capacity limit.
 - If the pools or storage systems have different capacity limits, enter 0 in the **Capacity Limit (%)** field and press the Backspace key.
- 6. Save your changes.

Capacity limit metrics

Monitor how much capacity is available for storing data when you set a capacity limit for your storage systems and pools.

Capacity Limit (%)

As an Administrator, you can set a **Capacity Limit (%)** for block storage systems, or pools, or both. In this scenario, you want to keep 20% of your usable capacity in reserve so you set the capacity limit to 80%.

Example: Administrator Sets Capacity Limit to 80% 100 GiB Capacity 100 GiB Capacity Used Capacity = 40 GiB Available Capacity = 60 GiB

To see the GiB or percentage value for the capacity limit, right-click any column heading on the **Block Storage Systems** or the **Pools** pages and add them.


To monitor how much capacity is available in relation to the capacity limit that you set, the used and available capacity values are adjusted and are shown in the following columns:

- Adjusted Used Capacity (%)
- · Capacity-to-Limit (GiB)

To see the values for **Adjusted Used Capacity (%)**, right-click any column heading on the **Block Storage Systems** page or **Pools** page and add it.

Adjusted Used Capacity (%)

To know how much capacity can be used without exceeding the capacity limit, monitor the value for **Adjusted Used Capacity (%)**.



The formula for calculating **Adjusted Used Capacity (%)** is (Used Capacity in GiB/Capacity Limit in GiB)*100. Using the values that are shown in the illustration above, the value is (40 GiB/80 GiB)* 100 or 50%. So, in this scenario, you can use 30% or 40 GiB of the usable capacity of the resource before you reach the capacity limit.

If the used capacity exceeds the capacity limit, the value for **Adjusted Used Capacity (%)** is over 100%.

Capacity-to-Limit (GiB)

Use the value for **Capacity-to-Limit (GiB)** to monitor how much available capacity is left before the capacity limit is reached.



The formula for calculating **Capacity-to-Limit (GiB)** is (Capacity Limit in GiB - Used Capacity in GiB). Using the values that are shown in the illustration above, the value is (80 GiB - 40 GiB) or 40 GiB.

If the capacity limit is exceeded, a negative value is shown for Capacity-to-Limit (GiB).

Viewing capacity information

View the capacity and space usage information of resources such as block storage systems, volumes, and pools. You can view similar information for filesets, file systems, and file system pools as internal resources of file storage systems. You can view information for containers as internal resources of object storage systems. Additionally, you can view capacity and space usage information for resources that violated a specified capacity range and generated an alert.

Before you view the capacity and space usage of resources, ensure that capacity data was collected for those resources during the time ranges that you want to analyze. Data collection is usually scheduled when resources are added for monitoring by IBM Spectrum Control.

You can view the following capacity information:

- "Capacity charts for block storage" on page 353
- "Capacity charts for file storage" on page 354
- "Capacity charts for object storage" on page 355
- "Capacity charts for tiers" on page 355

Capacity charts for block storage

Table 207. View capacity charts for block storage	
Actions	Navigation
View capacity charts for block storage systems.	From the Storage menu, click Block Storage Systems. Right-click one or more storage systems and click View Capacity.
View the capacity charts for all block storage pools.	From the Storage menu, click Pools . Right-click one or more pools and click View Capacity .
View the capacity chart for the pools in a storage system.	From the Storage menu, click Block Storage Systems.
	Right-click the storage system and click View Details.
	3. In the navigation menu, click Pools .
	4. Right-click one or more pools and click View Capacity .
View the capacity chart for all of the compressed, thin-provisioned, and Easy Tier volumes.	From the Storage menu, click Volumes . Right-click one or more volumes that are compressed, or thin-provisioned, or enabled for Easy Tier, and then click View Capacity .

Table 207. View capacity charts for block storage (continued)		
Actions	Navigation	
View the capacity chart for all of the compressed, thin-provisioned, and Easy Tier volumes in the storage system.	From the Storage menu, click Block Storage Systems.	
	Right-click the storage system and click View Details .	
	3. In the navigation menu, click Volumes .	
	4. Right-click one or more volumes that are compressed, or thin-provisioned, or enabled for Easy Tier, and then click View Capacity .	

Capacity charts for file storage

Table 208. View capacity charts for file storage		
Actions	Navigation	
View capacity charts for all of the filesets in a file storage system.	From the Storage menu, click File Storage Systems.	
	Right-click the storage system and click View Details.	
	3. In the navigation menu, click Filesets .	
	Right-click one or more filesets and click View Capacity.	
View the capacity charts for all of the file systems in a file storage system.	From the Storage menu, click File Storage Systems.	
	Right-click the storage system and click View Details .	
	3. In the navigation menu, click File Systems .	
	4. Right-click one or more file systems and click View Capacity.	
View the capacity charts for all of the file system pools in a file storage system.	From the Storage menu, click File Storage Systems.	
	Right-click the storage system and click View Details .	
	In the navigation menu, click File System Pools .	
	4. Right-click one or more file system pools and click View Capacity .	

Capacity charts for object storage

Table 209. View capacity charts for object storage		
Actions	Navigation	
View the capacity charts for all of the containers in an object storage system.	From the Storage menu, click Object Storage Systems.	
	Right-click the storage system and click View Details .	
	3. In the navigation menu, click Containers .	
	Right-click one or more containers and click View Capacity.	

Capacity charts for tiers

Table 210. View capacity charts for tiers		
Actions	Navigation	
View the capacity charts for tiers.	From the Groups menu, click Tiers . Right-click one or more tiers and click View Capacity .	

Related reference

Capacity metrics for block storage systems

To review trends in capacity and space usage for storage, you add metrics to capacity charts. You use the charts to detect capacity shortages and space usage trends.

Capacity metrics for file storage systems

To review trends in capacity and space usage for file storage systems, you add metrics to capacity charts. You use the charts for filesets, file systems, and file system pools to detect capacity shortages and space usage trends.

Capacity metrics for object storage systems

To review trends in capacity and space usage for object storage systems, you add metrics to capacity charts. Use the charts for containers to detect capacity shortages and space usage trends for the containers in your object storage systems.

Capacity metrics for tiers

Review trends in the capacity and space usage for tiers. You use the charts to detect capacity shortages and space usage for the tiers in your storage environment.

How capacity information is displayed

The capacity view is displayed when you view the capacity of resources, such as block storage systems or pools, including the capacity of resources in alert violations. Information in the capacity view is organized into a chart or table and a legend.

Capacity chart

The top section of the capacity view shows information about the selected resources, such as block storage systems or pools. You can view this information in the following formats:



The chart shows a visual representation of how the capacity of a resource trends over time. Each line on the chart represents a metric and a resource. For example, if you select two metrics and three resources, six lines are shown on the chart. The y-axis shows the unit of measurement for a metric. If more metrics were selected with a different unit of measurement, an extra y-axis is shown on the right side of the chart window.

Table

The table shows capacity information that is formatted into rows and columns. Each row represents a resource and a time stamp from the chart; each column represents a metric from the chart. For example, if two metrics and three resources are displayed on a chart, and each of the lines on the chart have 10 data points, 30 rows and 2 columns are shown in the table.

To view other metrics and asset information for a resource, right-click anywhere in the header row of the table and select extra columns. The type of resource determines the metrics and information that is available.

Capacity chart legend

The bottom section of the capacity view shows more information about the resources in the chart. The information is formatted into rows and columns. Each row represents a resource that was selected for the view. Each column provides capacity and space usage information about a resource.

You can also view information about the resources that are related to the resources in the legend. For example, in the capacity view for block storage systems, if a SAN Volume Controller is listed in the legend, right-click it to view the capacity of its internal resources, such as pools.

Tip:

- If you hide a resource in the capacity chart, the row for that resource remains visible in the chart legend.
- In the chart legend, information about a resource represents the state of that resource during the selected time range. Specifically, this information does not represent the current state of a resource, but instead shows the final state of the resource during the selected time range.

For example, if the current date is January 1, and you view a pool with the time range set to December 1 to December 7, the legend shows the capacity of that pool on December 7. If the capacity of the pool was changed between December 7 and January 1, this historical capacity is different from the current capacity of the pool on January 1.

Related reference

Capacity metrics

Use IBM Spectrum Control to collect and view capacity metrics about the storage systems in your environment.

Controls for capacity views

Each capacity view includes controls for customizing how information is displayed. The type of view and the resources that you are viewing determine which controls are available.



View chart

View capacity information in a chart format. The chart shows a visual representation of how resource capacity trends over time. Each line on the chart represents a metric and a resource. The y-axis shows the unit of measurement for a metric. If more metrics were selected with a different unit of measurement, an extra y-axis is shown on the right side of the chart window.

Hover the mouse pointer over points on a line to view a snapshot of capacity information at a specific time.



View table

View capacity information in a table format. Each row represents a resource and a time stamp. Each column represents a metric. You can view other metrics and information for a resource by right-clicking anywhere in the header row for the table and selecting more columns. The type of resource determines the metrics and information that are available.

For information how to filter, sort, and customize the columns in a table, see <u>"Customizing lists" on page 454.</u>



Metrics

Add metrics to the capacity chart. On the **Select Chart Metrics** dialog, the metrics that are available depend on the type of resource that is being shown in the chart. The number next to the name of the metrics type represents the number of metrics that are currently selected.

You can select multiple metrics at the same time, but you cannot include more than two unit types in the same view. For example, if you select metrics that use % and ops/s as units of measurement, you cannot select more metrics that use different units of measurement such as KiB/op or MiB/s.

For a list of metrics that you can view for resources, see "Capacity metrics" on page 367.



Hide and show resources

When the chart includes multiple resources, you can click the icon next to a resource to show only the line for that resource. Each icon is shown in a different color to match the color of line for the resource.

You can also show and hide resources in the chart by selecting resources in the chart legend. To select multiple resources at the same time, press Shift or Ctrl and click those resources. Press Shift and click to select consecutive rows in the legend; press Ctrl and click to select non-consecutive rows.

You can show up to 10 resources in a chart at the same time.

Specify a time range

The time range of the capacity information is shown below the chart. You can change this range to display information for different times when data was collected. When you first access the capacity view, the default time range is the last month.

Ensure that capacity data was collected during the time range that you select. If data was not collected during the time range, the chart and table are blank. If data collection was interrupted during the time range, the chart and table show gaps for the time increments when data was not collected.

For example, if you select a time range for the last 7 days, but data was not collected on days 4 and 5, the lines in the chart do not show data for days 4 and 5.



Export information about the chart to a file

Export information on a capacity view to a CSV file.



Open the capacity view in a separate web browser window

Open a duplicate of the current capacity view in a separate web browser window. You can change the information that is displayed in this separate window while retaining the original capacity view for comparison.

Actions for managing the resources in the chart legend

The chart legend in bottom section of the view shows more information about the selected resources. This information is organized into rows and columns, where each row represents a resource.

When you select one or more resources in the legend, the following actions are available in the **Actions** menu:

View Properties

View key details about a resource, including asset, status, configuration, and capacity information.

View resource Capacity

View the capacity of resources that are internal or related to a resource in the chart legend. For example, when you view the capacity of a SAN Volume Controller, you can right-click it and view the capacity of its internal resources, such as pools.

Information about an internal or related resource is shown in a separate web browser window. This window uses the same time range as the capacity view in the main window of the GUI.

Resources in the capacity chart legend

In the capacity view, capacity metrics and related information for resources are shown in a chart and in the chart legend. For resources in the chart legend, such as block storage systems, you can open separate capacity views for their internal resources, such as pools or volumes. The separate capacity views use the same time range as the capacity view in the main window of the GUI.

For example, when you view the capacity of a SAN Volume Controller, you can right-click it in the chart legend to open a capacity view for its internal resources, such as pools or volumes.

You can view capacity metrics for the following resources:

- Block storage systems
- Internal resources of block storage systems: volumes and pools
- Internal resources of file storage systems: filesets, file systems, and file system pools
- Internal resources of object storage systems: containers
- Tiers


The following table summarizes the resources and the internal resources for which you can view capacity metrics.

Table 211. Internal resources for which you can view capacity metrics		
Resource	Internal resources	
Block storage system	VolumesPools	
File storage system	FilesetsFile systemsFile system pools	
Object storage system	Containers	

Investigating capacity trends for block storage systems

Use the capacity charts and the information that is shown in the tables for the block storage systems to check which storage systems have the highest growth rates and which storage systems might need more capacity.

- 1. From the **Storage** menu, click **Block Storage Systems**.
- 2. Click **Capacity** and select the storage systems and the capacity metrics that you want to review.




Tip: To check which pools consume the most storage, right-click the storage systems and click **Capacity**.

Investigating capacity trends for block storage pools

Use the capacity charts and the information that is shown in the tables for the block storage pools to check which storage pools have the highest growth rates, which storage pools might need more capacity, and which storage pools are at risk because too much capacity is committed to the thin-provisioned volumes.

- From the Storage menu, click Block Storage Systems.
 Alternatively, you can click Storage > Pools.
- 2. Right-click the storage system and click **View Details**.
- 3. In the navigation pane, click **Pools**.
- 4. Click Capacity and select the pools and the capacity metrics that you want to review.




Tip: To investigate the pools with the highest shortfall rates, right-click the pools in the table and click **Volume Capacity**.

Investigating capacity trends for volumes

Use the capacity charts and the information that is shown in the tables to check the rate at which the thinprovisioned volumes consume capacity, the compression savings for the volumes that were converted to compressed volumes, and the changes in the distribution of Easy Tier volumes across the SSD, Enterprise HDD, and Nearline HDD drives.

- 1. From the Storage menu, click Block Storage Systems.
- 2. Right-click the storage system and click View Details.
- 3. In the navigation pane, click Volumes.
- 4. Click **Capacity** and select the volumes and the capacity metrics that you want to review.



Tip: To check the distribution of Easy Tier volumes, click **Select Chart Metrics** and add Tier metrics to the chart.

Investigating capacity trends for file systems

Use the capacity charts and the information that is shown in the tables for the file systems to check which file systems have the highest growth rates and which file systems might require additional capacity.

- 1. From the **Storage** menu, click **File Storage Systems**.
- 2. Right-click the storage system and click View Details.
- 3. In the navigation menu, click **File Systems**.
- 4. Click Capacity and select the file systems and the capacity metrics that you want to review.




Tip: To check which pools consume the most storage, right-click the file systems and click **File System Pool Capacity**.

Investigating capacity trends for file system pools

Use the capacity charts and the information that is shown in the tables for the file system pools in your storage systems to check which storage pools have the highest growth rates so that you can plan future storage requirements.


- 1. From the Storage menu, click File Storage Systems.
- 2. Right-click the storage system and click View Details.
- 3. In the navigation menu, click File System Pools.
- 4. Click **Capacity** and select the pools and the capacity metrics that you want to review.



Investigating capacity trends for filesets

Use the capacity charts to check which filesets have the highest growth rates in your file storage systems.

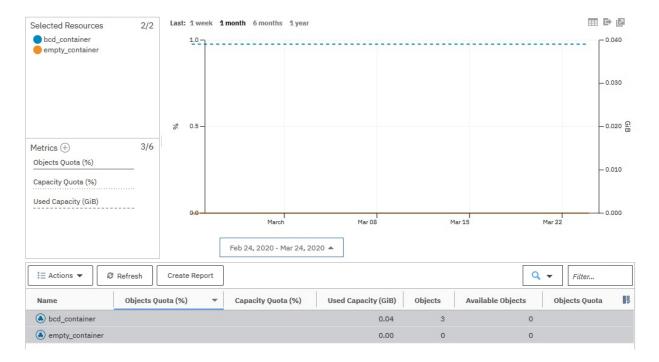
- 1. From the **Storage** menu, click **File Storage Systems**.
- 2. Right-click the storage system and click View Details.
- 3. In the navigation menu, click Filesets.
- 4. Click **Capacity** and select the filesets that you want to review.



Investigating capacity trends for containers

Use the capacity charts to check which containers have the highest growth rates in your object storage systems.

- 1. From the **Storage** menu, click **Object Storage Systems**.
- 2. Right-click the storage system and click **View Details**.
- 3. In the navigation menu, click **Containers**.
- 4. Click **Capacity** and select the objects that you want to review.



Investigating capacity trends for tiers

Use the capacity charts and the information that is shown in the tables to check which tiers have the highest growth rates and which tiers might require more capacity.

- 1. From the **Groups** menu, click **Tiers**.
- 2. Click Capacity.

Creating bookmarks for URLs of capacity views

You can create a bookmark for the URL of a capacity view. You can also open a duplicate of the capacity view to change and compare views.

The bookmark stores details of the capacity view, such as the resources, metrics, interval, and other details. It also stores the time period that is specified in the time selector at the top of the chart.

To create a bookmark for a capacity view, follow these steps:

- 1. Open the capacity view for which you want to create a bookmark.
- 2. Right-click and click the relevant option for creating a bookmark.

Tip: To open a duplicate of the current capacity view, click the icon for opening a new browser window. You can change the information that is displayed in the new window while retaining the original capacity view for comparison.

Viewing the capacity of external storage

View the used capacity and capacity of external pool storage that is used by file systems in IBM Spectrum Scale. External pools can include storage that is provided by IBM Cloud Object Storage, Amazon Simple Storage Service (S3), OpenStack Swift, IBM Spectrum Archive, IBM Spectrum Protect, and other storage providers.


For all external storage that is being used by IBM Spectrum Scale file systems, you can complete the following tasks:

- View used capacity to understand how much data is being migrated from those file systems to external storage.
- Charge back storage costs to the departments within your organization.
- Gauge shortfall in case you need to recall data from external pools.

For external storage that is provided by IBM Cloud Object Storage, you can also view capacity information, including the percentage of space (active and inactive) that is being used. Use this information to complete the following tasks:

- View how storage space is distributed between internal and external storage tiers.
- Identify the external pools that are running out of space.
- Determine when you need to reconcile files between an IBM Spectrum Scale and an external storage tier. For example, when you remove files from a file system, no policy is available to automatically remove the cloud objects that are created. This situation might cause orphan objects (inactive data) on the cloud. To clean up inactive data and keep the cloud in sync with the file system, you can reconcile files.

The following figure illustrates the relationship between IBM Spectrum Scale and the IBM Cloud Object Storage system that provides storage for external pools:



To view the used capacity and capacity of external storage, complete the following steps:

- 1. From the Storage menu, click File Storage Systems.
- 2. Right-click a IBM Spectrum Scale storage system and click View Details.
- 3. In the navigation pane, click **Pools**.
- 4. View the following columns for information about the used capacity and capacity of external pools:

Capacity (%)

The percentage of used capacity in the file system pool. Hover over the bar for a specific pool to view statistics about its used and available capacity.

For external pools that are connected to cloud services or storage providers, only the amount of used capacity is shown. Cloud services and other storage providers might include IBM Spectrum Protect, IBM Spectrum Archive, Amazon Simple Storage Service (S3), or OpenStack Swift.

For external pools that are connected to vaults in IBM Cloud Object Storage, the percentage of used capacity includes active and inactive data. Space values are determined from the usable vault capacity that is associated with the cloud account and the hard quota value (if a hard quota is configured). Hover over the bar for a specific pool to view statistics about its used active space, used inactive space, and available capacity.

Active and inactive data in external pools

Active data is data that has corresponding stub files on the GPFS file system. Inactive data is data that exists in the external pool but is not accessible in the GPFS file system. Data is classified as inactive under these conditions:

- When you migrate a file off a GPFS file system, its data is still visible and accessible through its stub file. When the stub file is deleted, the data in the external pool is no longer accessible and is considered inactive.
- When you modify a migrated file, it's recalled from the external pool and modified locally. The previous data for the file remains in the external pool but is no longer accessible through the GPFS file system and is considered inactive.

Tip: To keep the files in sync between an external pool and a GPFS file system and clean up inactive data, you can run the IBM Spectrum Scale mmcloudgateway files reconcile command. For more information, see Reconciling files between IBM Spectrum Scale file system and cloud storage tier.

External Used Capacity (GiB)

The amount of used capacity on an external pool. This value includes active data and inactive data. Active data is data that has corresponding stub files on the GPFS file system. Inactive data is data that exists in the external pool but is not accessible in the GPFS file system.

Available for: External pools that are provided by IBM Cloud Object Storage.

Inactive Used Capacity

The amount of inactive used capacity on an external pool. Inactive data is data that exists in the external pool but is not accessible in the GPFS file system. Available for external pools that are provided by Cloud Object Storage only.

Available for: External pools that are provided by IBM Cloud Object Storage.

Used Capacity

The amount of capacity in the file system pool that is being used. For external pools, used capacity includes active data only. Active data is data that has corresponding stub files on the GPFS file system.

Available for: All storage systems.

Related information

Introduction to Transparent Cloud Tiering for IBM Spectrum Scale

Identifying shortfall before data is recalled from external storage

View the percentage of migrated data in external pools that does not fit into the available capacity on an IBM Spectrum Scale file system.

This percentage, called *shortfall*, represents the relative risk of not having enough internal file system space if you recall all data from the associated external pools. The higher the percentage, the higher the amount of migrated data that does not fit.

To view the potential shortfall for a file system, complete the following steps:

- 1. From the Storage menu, click File Storage Systems.
- 2. Right-click a IBM Spectrum Scale storage system and click View Details.
- 3. Under Internal Resources, click File Systems.
- 4. Locate the file system that you want to check and view the value in the Shortfall column.

To calculate shortfall, the following formula is used:

```
[(External Pool Migrated Data - Available Capacity) ÷ External Pool Migrated Data] × 100
```

For example, if the available capacity on a file system is 200 GiB, but the used capacity of migrated data on external pools is 500 GiB, the shortfall percentage is 60% (300 GiB). In this case, consider adding capacity to the file system before recalling data.

```
[(500 - 200) \div 500] \times 100 = 60\%
```

If the file system has sufficient available capacity to contain all the migrated data in external pools, no shortfall exists. For example, if the available capacity on a file system is 500 GiB, and the used capacity in external pools is only 200 GiB, the shortfall is 0%.

Tip: Only the space that is used by migrated data in external pools is included in the calculation for shortfall. Pre-migrated data is not included because it exists in both the internal and external pools.

Viewing capacity alerts and violations

View the alerts that were triggered when the capacity of a resource changes and reaches a threshold. For example, you can view alerts that are generated when the measured value of a capacity metric for a pool meets the conditions for generating an alert.

To view alert violations for the capacity of a resource, you must collect capacity data and define capacity alerts for that resource or for the alert policy that manages the resource.

- 1. For block storage systems, go to **Storage** > **Block Storage Systems**. For file storage systems, go to **Storage** > **File Storage Systems**.
- 2. Right-click a storage system and click View Details.
- 3. In the navigation menu, click **Alerts**.
- 4. Right-click the resource alert that you want to view, such as an alert that was defined for a capacity metric for the resource, and click **View Alert**.

Related concepts

Collecting data

Determining the data that you want to gather about storage resources is critical to helping you implement a storage management strategy. IBM Spectrum Control provides two different jobs for collecting data about resources.

Alerting

Specify conditions that trigger alerts and the actions to take when those alerts are triggered, such as notify an email address. Use alert policies to define those alert conditions and notification settings for a group of resources.

Related tasks

Defining alert definitions for general attributes and capacity changes

You can define alerts that are triggered when the attributes or capacity of a resource changes. Attributes represent the key properties and configuration of a resource, such as status, versions, removals, discoveries, and data collection status. Capacity represents storage statistics such as available capacity, used capacity, drive capacity, reserved capacity, and more.

Capacity metrics

Use IBM Spectrum Control to collect and view capacity metrics about the storage systems in your environment.

Capacity metrics for block storage systems

To review trends in capacity and space usage for storage, you add metrics to capacity charts. You use the charts to detect capacity shortages and space usage trends.

Alphabetical lists of the capacity and space usage metrics that you can add to charts are provided in the following sections:

- "Storage system capacity metrics" on page 367
- "Pool capacity metrics" on page 372
- "Volume capacity metrics" on page 380

Tip: IBM Spectrum Control displays capacity values in base 2 (GiB), while the XIV management GUI displays capacity values in base 10 (GB, TB). Even though different units of measurement are used, the storage values are equivalent. For more information about units of measurement, see "Units of measurement for storage data" on page 79.

Storage system capacity metrics

To detect capacity shortages and investigate space usage trends, you can add the following metrics to the capacity chart for storage systems:

Adjusted Used Capacity (%)

The amount of capacity that can be used without exceeding the capacity limit.

Example: Adjusted Used Capacity

The formula for calculating **Adjusted Used Capacity (%)** is (Used Capacity in GiB/Capacity Limit in GiB)*100. For example, if the capacity is 100 GiB, the used capacity is 40 GiB, and the capacity limit is 80% or 80 GiB, then the value for **Adjusted Used Capacity (%)** is (40 GiB/80 GiB)* 100 or 50%. So, in this example, you can use 30% or 40 GiB of the usable capacity of the resource before you reach the capacity limit.

If the used capacity exceeds the capacity limit, the value for **Adjusted Used Capacity (%)** is over 100%.

To add the **Adjusted Used Capacity (%)** column, right-click any column heading on the **Block Storage Systems** page.

See these related values for more information Capacity Limit (%), and Capacity-to-Limit (GiB).

Available Capacity (GiB)

(Previously known as Available Pool Space) The total amount of the space in the pools that is not used by the volumes in the pools. To calculate available capacity, the following formula is used:

```
(pool capacity - used capacity)
```

For XIV systems, pool capacity is the physical capacity of the pools and does not include the provisioned capacity of the pools.

Available for: All storage systems.

Available Volume Capacity (GiB)

(Previously known as Effective Unallocated Volume Space) The total amount of remaining space that can be used by the volumes in the pools. The following formula is used to calculate this value:

```
[Provisioned Capacity - Used Capacity]
```

The capacity that is used by thin-provisioned volumes is typically less than their provisioned capacity. Therefore, the available capacity represents the difference between the provisioned capacity and the used capacity for all the volumes in the pools. For Hitachi VSP non-thin provisioned pool capacity, the available capacity is always zero.

Available for: All storage systems.

Capacity (GiB)

(Previously known as Pool Capacity) The total amount of storage space in the pools. For XIV systems and IBM Spectrum Accelerate, capacity represents the physical ("hard") capacity of the pool, not the provisioned ("soft") capacity. Pools that are allocated from other pools are not included in the total pool space.

Available for: All storage systems.

Capacity Limit (%) and Capacity Limit (GiB)

The limit that was set on the capacity that is used by your storage systems. For example, the policy of your company is to keep 20% of the usable capacity of your storage systems in reserve. So, you log into the GUI as Administrator and set the capacity limit to 80%.

Example: Administrator Sets Capacity Limit to 80%



1. Set the capacity limit



Click the illustration above to find out how to set capacity limits.

The GiB value for the capacity limit for the storage system is calculated when you set the value for the **Capacity Limit (%)**.


To add the **Capacity Limit (%)** and the **Capacity Limit (GiB)** columns, right-click any column heading on the **Block Storage Systems** page.

See these related values for more information **Adjusted Used Capacity (%)** and **Capacity-to-Limit** (GiB).

Capacity-to-Limit (GiB)

The amount of capacity that is available before the capacity limit is reached.

Example: Capacity-to-Limit



The formula for calculating **Capacity-to-Limit (GiB)** is (Capacity Limit in GiB - Used Capacity in GiB). For example, if the capacity limit is 80% or 80 GiB and the used capacity is 40 GiB, then the value for **Capacity-to-Limit (GiB)** is (80 GiB - 40 GiB) or 40 GiB.

See these related values for more information Capacity Limit (%) and Adjusted Used Capacity (%).

Compression Savings (%)

The estimated amount and percentage of capacity that is saved by using data compression, across all pools on the storage system. The percentage is calculated across all compressed volumes in the pools and does not include the capacity of non-compressed volumes.

For storage systems with drives that use inline data compression technology, the Compression Savings does not include the capacity savings that are achieved at the drive level.

The following formula is used to calculate the amount of storage space that is saved:

```
written capacity - compressed size
```

The following formula is used to calculate the percentage of capacity that is saved:

```
((written capacity - compressed size) \div written capacity) \times 100
```

For example, the written capacity, which is the amount of data that is written to the volumes before compression, is 40 GiB. The compressed size, which reflects the size of compressed data that is written to disk, is just 10 GiB. Therefore, the compression savings percentage across all compressed volumes is 75%.

Available for: FlashSystem A9000 and FlashSystem A9000R, IBM Spectrum Accelerate, XIV storage systems with firmware version 11.6 or later, and resources that run IBM Spectrum Virtualize.

For FlashSystem A9000 and FlashSystem A9000R, all volumes in the pools are compressed.

Deduplication Savings (%)

The estimated amount and percentage of capacity that is saved by using data deduplication, across all data reduction pools on the storage system. The percentage is calculated across all deduplicated volumes in the pools and does not include the capacity of volumes that are not deduplicated.

The following formula is used to calculate the amount of storage space that is saved:

```
written capacity - deduplicated size
```

The following formula is used to calculate the percentage of capacity that is saved:

```
((written capacity - deduplicated size) ÷ written capacity) × 100
```

For example, the written capacity, which is the amount of data that is written to the volumes before deduplication, is 40 GiB. The deduplicated size, which reflects the size of deduplicated data that is written to disk, is just 10 GB. Therefore, data deduplication reduced the size of the data that is written by 75%.

Available for: FlashSystem A9000, FlashSystem A9000R, and resources that run IBM Spectrum Virtualize version 8.1.3 or later.

Drive Compression Savings (%)

The amount and percentage of capacity that is saved with drives that use inline data compression technology. The percentage is calculated across all compressed drives in the pools.

The amount of storage space that is saved is the sum of drive compression savings.

The following formula is used to calculate the percentage of capacity that is saved:

```
((written capacity - compressed size) ÷ written capacity) × 100
```

Available for: Storage systems that contain IBM FlashCore® Modules with hardware compression.

Mapped Capacity (GiB)

(Previously known as Assigned Volume Space) The total volume space in the storage system that is mapped or assigned to host systems, including child pool capacity.

Available for: All storage systems.

Overprovisioned Capacity (GiB)

(Previously known as Unallocatable Volume Space) The capacity that cannot be used by volumes because the physical capacity of the pools cannot meet the demands for provisioned capacity. The following formula is used to calculate this value:

```
[Provisioned Capacity - Capacity]
```

Available for: All storage systems.

Shortfall (%)

The percentage of space that is over committed to the pools with thin-provisioned volumes. For example, you commit 100 GiB of space to a thin-provisioned volume in a pool with a capacity of 50 GiB. As the space is used by the thin-provisioned volume in increments of 10 GiB, the space available for allocation decreases and the shortfall in capacity becomes more acute.

To calculate the shortfall, the following formula is used:

```
[(overprovisioned capacity ÷ committed but available capacity) × 100]
```

A shortfall occurs when you commit more space to the volumes in the pools than is physically available to the pools. If the physical space available to the pools is less than the committed

provisioned capacity, then the pools do not have enough space to fulfill the commitment to the provisioned capacity.

For example, the physical capacity of the pools is 70 GiB, but 150 GiB of provisioned capacity was committed to the thin-provisioned volumes. If the volumes are using 50 GiB, 100 GiB is committed to those volumes (150 GiB – 50 GiB) with only 20 GiB of available pool space (70 GiB – 50 GiB). Because only 20 GiB of the pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% [($80 \text{ GiB} \div 100 \text{ GiB} \times 100$].

Available for: DS8000, FlashSystem storage systems, Dell EMC VMAX, VNX, and VNXe storage systems, Hitachi Virtual Storage Platform, SAN Volume Controller, Storwize family storage systems that are configured with block storage, XIV systems, and IBM Spectrum Accelerate storage systems.

Provisioned Capacity (%)

(Previously known as Virtual Allocation) The percentage of the physical capacity that is committed to the provisioned capacity of the volumes in the pools. If the value exceeds 100%, the physical capacity doesn't meet the demands for provisioned capacity.

To calculate provisioned capacity percentage, the following formula is used:

```
[(provisioned capacity ÷ pool capacity) × 100]
```

For example, if the provisioned capacity percentage is 200% for a storage pool with a physical capacity of 15 GiB, then the provisioned capacity that is committed to the volumes in the pools is 30 GiB. Twice as much space is committed to the pools than is physically available to the pools. If the provisioned capacity percentage is 100% and the physical capacity is 15 GiB, then the provisioned capacity that is committed to the pools is 15 GiB. The total physical capacity that is available to the pools is used by the volumes in the pools.

A provisioned capacity percentage that is higher than 100% is considered to be aggressive because insufficient physical capacity is available to the pools to satisfy the allocation of the committed space to the compressed and thin-provisioned volumes in the pools. In such cases, you can check the **Shortfall (%)** value to determine how critical the shortage of space is for the storage system pools.

Available for: All storage systems.

Provisioned Capacity (GiB)

(Previously known as Total Volume Capacity) The total amount of storage space that can be made available to the standard- and thin-provisioned volumes in the pools. If the pool is a parent pool, it also includes the storage space that can be made available to the volumes in the child pools.

Available for: All storage systems.

Total Capacity Savings (%)

(Previously known as Total Data Reduction Savings) The estimated amount and percentage of capacity that is saved by using data deduplication, pool compression, thin provisioning, and drive compression, across all volumes in the pool.

The following formula is used to calculate the amount of storage space that is saved:

```
[Provisioned Capacity - Used Capacity]
```

The following formula is used to calculate the percentage of capacity that is saved:

```
((Provisioned Capacity - Used Capacity) ÷ Provisioned Capacity) × 100
```

Available for: FlashSystem A9000 and FlashSystem A9000R, IBM Spectrum Accelerate, XIV storage systems with firmware version 11.6 or later, and resources that run IBM Spectrum Virtualize.


Unmapped Capacity (GiB)

(Previously known as Unassigned Volume Space) The total amount of space in the volumes that are not assigned to hosts.

Available for: All storage systems.

Used Capacity (%)

(Previously known as Physical Allocation)



The percentage of physical capacity in the pools that is used by the standard-provisioned volumes, the thin-provisioned volumes, and the volumes in child pools. Check the value for used capacity percentage to see:

- Whether the physical capacity of the pools is fully allocated. That is, the value for used capacity is 100%.
- Whether you have sufficient capacity to provision new volumes with storage
- Whether you have sufficient capacity to allocate to the compressed and thin-provisioned volumes in the pools

Available for: All storage systems.

Used Capacity (GiB)

(Previously known as Allocated Space) The amount of space that is used by the standard- and thin-provisioned volumes in the pools. If the pool is a parent pool, the amount of space that is used by the volumes in the child pools is also calculated.

The capacity that is used by for thin-provisioned volumes is less than their provisioned capacity, which is shown in the **Provisioned Capacity (GiB)** column. If a pool doesn't have thin-provisioned volumes, the value for used capacity is the same as the value for provisioned capacity.

Available for: All storage systems.

Pool capacity metrics

If sufficient data is collected, you can view charts that compare the capacity, used capacity, and available capacity of the pools in your data center.

In the **Zero Capacity** column on the **Pools** page, you can see the date, based on the storage usage trends for the pool, when the pool will run out of available capacity.

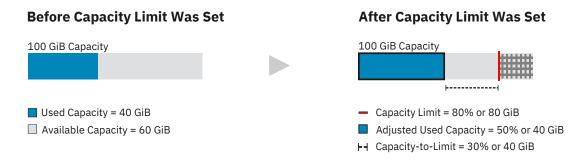
Zero Capacity: The capacity information that is collected over 180 days is analyzed to determine, based on historical storage consumption, when the pools will run out of capacity. The pools that have already run out of capacity are marked as depleted. For the other pools, a date is provided so that you know when the pools are projected to run out of capacity. If sufficient information isn't collected to analyze the storage usage of the pool, None is shown as the value for zero capacity. If a capacity limit is set for the pool, the date shown in the **Zero Capacity** column is the date when the available capacity based on the capacity limit will be depleted. For example, if the capacity limit for a 100 GiB pool is 80%, it is the date when the available capacity of the pool is less than 20 GiB. Depleted is shown in the column when the capacity limit is reached.

To detect capacity shortages and investigate trends in storage usage, you can add the following metrics to the capacity chart for pools:

Adjusted Used Capacity (%)

The amount of capacity that can be used without exceeding the capacity limit.

Example: Adjusted Used Capacity



The formula for calculating **Adjusted Used Capacity (%)** is (Used Capacity in GiB/Capacity Limit in GiB)*100. For example, if the capacity is 100 GiB, the used capacity is 40 GiB, and the capacity limit is 80% or 80 GiB, then the value for **Adjusted Used Capacity (%)** is (40 GiB/80 GiB)* 100 or 50%. So, in this example, you can use 30% or 40 GiB of the usable capacity of the resource before you reach the capacity limit.

If the used capacity exceeds the capacity limit, the value for **Adjusted Used Capacity (%)** is over 100%.

To add the Adjusted Used Capacity (%) column, right-click any column heading on the Pools page.

See these related values for more information Capacity Limit (%) and Capacity-to-Limit (GiB).

Available Capacity (GiB)

(Previously known as Available Pool Space) The amount of space that is available to create new volumes in the pool. If the pool is a parent pool, the amount of space that is used by the volumes in the child pools is also included.

Available for: All storage systems. For FlashSystem A9000 and FlashSystem A9000R, this value represents provisioned capacity rather than physical space.

Available Repository Capacity (GiB)

The available, unallocated storage space in the repository for Track Space-Efficient (TSE) thin-provisioning.

Available for: DS8000 thin-provisioned pools.

Available Soft Capacity (GiB)

The amount of virtual storage space that is available to allocate to volumes in a storage pool.

Available for: XIV systems, and IBM Spectrum Accelerate storage systems.

Available Written Capacity (GiB)

(Previously known as Effective Used Capacity) For devices with inline hardware compression, the effective available capacity is the total amount of the provisioned capacity in the pools that is not used by the volumes in the pools.

Capacity (GiB)

The total amount of storage space in the pool. For XIV systems and IBM Spectrum Accelerate, capacity represents the physical or ("hard") capacity of the pool, not the provisioned ("soft") capacity.

Available for: All storage systems.

Capacity Limit (%) and Capacity Limit (GiB)

The limit that was set on the capacity that is used by your pools. For example, the policy of your company is to keep 20% of the usable capacity of your pools in reserve. So, you log into the GUI as Administrator and set the capacity limit of your pools to 80%.

Example: Administrator Sets Capacity Limit to 80%



1. Set the capacity limit



Click the illustration above to find out how to set capacity limits.

The GiB value for the capacity limit for the pool is calculated when you set the value for the **Capacity Limit (%)**.

To add the **Capacity Limit (%)** and the **Capacity Limit (GiB)** columns, right-click any column heading on the **Pools** page.

See these related values for more information **Adjusted Used Capacity (%)** and **Capacity-to-Limit** (GiB).

Zero capacity: When you set the capacity limit for pools, the values shown for **Zero Capacity** are readjusted to take into account the capacity limit of the pool. The date will represent when the capacity limit of the pool is reached. If the pool has already reached the capacity limit, Depleted is shown. None is shown when a trend in storage consumption can't be detected because the pool's storage isn't being consumed or because not enough data was collected to predict storage consumption.

Capacity-to-Limit (GiB)

The amount of capacity that is available before the capacity limit is reached.

Example: Capacity-to-Limit 100 GiB Capacity — Capacity Limit = 80% or 80 GiB Adjusted Used Capacity = 50% or 40 GiB H Capacity-to-Limit = 30% or 40 GiB

The formula for calculating **Capacity-to-Limit (GiB)** is (Capacity Limit in GiB - Used Capacity in GiB). For example, if the capacity limit is 80% or 80 GiB and the used capacity is 40 GiB, then the value for **Capacity-to-Limit (GiB)** is (80 GiB - 40 GiB) or 40 GiB.

See these related values for more information Capacity Limit (%) and Adjusted Used Capacity (%).

Compression Savings (%)

The estimated amount and percentage of capacity that is saved by using data compression. The percentage is calculated across all compressed volumes in the pool and does not include the capacity of non-compressed volumes.

For storage systems with drives that use inline data compression technology, the Compression Savings does not include the capacity savings that are achieved at the drive level.

The following formula is used to calculate the amount of storage space that is saved:

```
written capacity - compressed size
```

The following formula is used to calculate the percentage of capacity that is saved:

```
((written capacity - compressed size) ÷ written capacity) × 100
```

For example, the written capacity, which is the amount of data that is written to the volumes before compression, is 40 GiB. The compressed size, which reflects the size of compressed data that is written to disk, is just 10 GiB. Therefore, the compression savings percentage across all compressed volumes is 75%.

Available for: FlashSystem A9000 and FlashSystem A9000R, IBM Spectrum Accelerate, XIV storage systems with firmware version 11.6 or later, and resources that run IBM Spectrum Virtualize.

For FlashSystem A9000 and FlashSystem A9000R, all volumes in the pools are compressed.

Deduplication Savings (%)

The estimated amount and percentage of capacity that is saved by using data deduplication. The percentage is calculated across all deduplicated volumes in the pool and does not include the capacity of volumes that are not deduplicated.

The following formula is used to calculate the amount of storage space that is saved:

```
written capacity - deduplicated size
```

The following formula is used to calculate the percentage of capacity that is saved:

```
((written capacity - deduplicated size) \div written capacity) \times 100
```

For example, the written capacity, which is the amount of data that is written to the volumes before deduplication, is 40 GiB. The deduplicated size, which reflects the size of deduplicated data that is written to disk, is just 10 GB. Therefore, data deduplication reduced the size of the data that is written by 75%.

Available for: Storage systems that run IBM Spectrum Virtualize version 8.1.3 or later.

Drive Compression Savings (%)

The amount and percentage of capacity that is saved with drives that use inline data compression technology. The percentage is calculated across all compressed drives in the pools.

The amount of storage space that is saved is the sum of drive compression savings.

The following formula is used to calculate the percentage of capacity that is saved:

```
((written capacity - compressed size) \div written capacity) \times 100
```

Available for: Storage systems that contain IBM FlashCore® Modules with hardware compression.

Enterprise HDD Available Capacity (GiB)

The amount of storage space that is available on the Enterprise hard disk drives that can be used by Easy Tier for re-tiering the volume extents in the pool.

Available for: DS8000 and storage systems that run IBM Spectrum Virtualize.

Enterprise HDD Capacity (GiB)

The total amount of storage space on the Enterprise hard disk drives that can be used by Easy Tier for re-tiering the volume extents in the pool.

Available for: DS8000 and storage systems that run IBM Spectrum Virtualize.

Mapped Capacity (GiB)

(Previously known as Assigned Volume Space) The total amount of space in the volumes that is assigned to hosts. For Hitachi VSP non-thin provisioning pool space, this value is the sum of assigned regular host-accessible volumes. Volumes that are used for thin-provisioning (pool volumes) are not included.

Available for: All storage systems.

Nearline HDD Available Capacity (GiB)

The amount of storage space that is available on the Nearline hard disk drives that can be used by Easy Tier for re-tiering the volume extents in the pool.

Available for: DS8000 and storage systems that run IBM Spectrum Virtualize.

Nearline HDD Capacity (GiB)

The total amount of storage space on the Nearline hard disk drives that can be used by Easy Tier for re-tiering the volume extents in the pool.

Available for: DS8000 and storage systems that run IBM Spectrum Virtualize.

Overprovisioned Capacity (GiB)

(Previously known as Unallocatable Volume Space) The capacity that cannot be used by volumes because the physical capacity of the pool cannot meet the demands for provisioned capacity. The following formula is used to calculate this value:

```
[Provisioned Capacity - Capacity]
```

In thin-provisioned environments, it is possible to over commit (over provision) storage in a pool by creating volumes with more provisioned capacity than can be physically allocated in the pool. This value represents the amount of volume capacity that cannot be allocated based on the current capacity of the pool. For Hitachi VSP non-thin provisioned pool capacity, this value is always zero.

Available for: All storage systems, except FlashSystem A9000, FlashSystem A9000R, XIV, and IBM Spectrum Accelerate.

Provisioned Capacity (%)

(Previously known as Virtual Allocation) The percentage of the physical capacity that is committed to the provisioned capacity of the volumes in the pool. If the value exceeds 100%, the physical capacity doesn't meet the demands for provisioned capacity. The following formula is used to calculate this value:

```
[(Provisioned Capacity ÷ Capacity) × 100]
```

This value is available for all pools.

For Hitachi VSP non-thin provisioned pool space, the following formula is used to calculate this value:

```
[(Used Capacity ÷ Capacity) × 100]
```

Example: If the provisioned capacity percentage is 200% for a total storage pool size of 15 GiB, then the provisioned capacity that is committed to the volumes in the pool is 30 GiB. This configuration means that twice as much capacity is committed than is physically contained in the pool. If the provisioned capacity percentage is 100% for the same pool, then the provisioned capacity that is committed to the pool is 15 GiB. This configuration means that all the physical capacity of the pool is already used by volumes.

A provisioned capacity percentage that is higher than 100% is considered aggressive because insufficient physical capacity is available in the pool to satisfy the maximum allocation for all the thin-provisioned volumes in the pool. In such cases, you can use the value for Shortfall (%) to estimate how critical the shortage of capacity is for a pool.

You can hover the mouse pointer over the percentage bar to view values for the provisioned capacity and capacity.

Available for: All storage systems.

Provisioned Capacity (GiB)

(Previously known as Total Volume Capacity) The total amount of storage capacity that can be made available to the standard- and thin-provisioned volumes in the pool. If the pool is a parent pool, it also includes the storage capacity that can be made available to the volumes in the child pools. For Hitachi VSP non-thin provisioned pool capacity, this value is the sum of the capacity of regular host-accessible volumes. Volumes that are used for thin-provisioning (pool volumes) are not included.

Available for: All storage systems.

Repository Capacity (GiB)

The total storage capacity of the repository for Track Space-Efficient (TSE) thin-provisioning.

Available for: DS8000 thin-provisioned pools.

Reserved Volume Capacity

(Previously known as Unused Space) The amount of pool capacity that is reserved but has not been used yet to store data on the thin-provisioned volume.

Available for: Resources that run IBM Spectrum Virtualize.

SCM Available Capacity (GiB)

The available capacity on Storage Class Memory (SCM) drives in the pool. Easy Tier can use these drives to retier the volume extents in the pool.

Available for: IBM Spectrum Virtualize systems, such as FlashSystem 9100, FlashSystem 7200, and Storwize family storage systems that are configured with block storage.

SCM Capacity (GiB)

The total capacity on Storage Class Memory (SCM) drives in the pool. Easy Tier can use these drives to retier the volume extents in the pool.

Available for: IBM Spectrum Virtualize systems, such as FlashSystem 9100, FlashSystem 7200, and Storwize family storage systems that are configured with block storage.

Shortfall (%)

The difference between the remaining unused volume capacity and the available capacity of the associated pool, expressed as a percentage of the remaining unused volume capacity. The shortfall represents the relative risk of running out of space for overallocated thin-provisioned volumes. If the pool has sufficient available capacity to satisfy the remaining unused volume capacity, no shortfall exists. As the remaining unused volume capacity grows, or as the available pool capacity decreases, the shortfall increases and the risk of running out of space becomes higher. If the available capacity of the pool is exhausted, the shortfall is 100% and any volumes that are not yet fully allocated have run out of space.

If the pool isn't thin-provisioned, the shortfall percentage equals zero. If shortfall percentage isn't calculated for the storage system, the field is left blank.

The following formula is used to calculate this value:

```
Overprovisioned Capacity ÷ Committed but Unused Capacity
```

You can use this percentage to determine when the amount of over-committed space in a pool is at a critically high level. Specifically, if the physical space in a pool is less than the committed provisioned capacity, then the pool does not have enough space to fulfill the commitment to provisioned capacity. This value represents the percentage of the committed provisioned capacity that is not available in a pool. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

View a short animation about about how Shortfall works.

Example: The remaining physical capacity of a pool is 70 GiB, but 150 GiB of provisioned capacity was committed to thin-provisioned volumes. If the volumes are using 50 GiB, then 100 GiB is still committed to the volumes (150 GiB – 50 GiB) with a shortfall of 30 GiB (70 GiB remaining pool space – 100 GiB remaining commitment of volume space to the volumes).

Because the volumes are overcommitted by 30 GiB based on the available capacity in the pool, the shortfall is 30% when the following calculation is used:

```
[(100 GiB unused volume capacity – 70 GiB remaining pool capacity)
÷ 100 GiB unused volume capacity] × 100
```

Available for: DS8000, Hitachi Virtual Storage Platform, and storage systems that run IBM Spectrum Virtualize.

For FlashSystem A9000 and FlashSystem A9000R, this value is not available.

Soft Capacity (GiB)

The amount of virtual storage space that is configured for the pool.

Available for: XIV systems and IBM Spectrum Accelerate storage systems.

Tier O Flash Available Capacity (GiB)

The amount of storage space that is available on the Tier 0 flash solid-state drives that can be used by Easy Tier for retiering the volume extents in the pool.

Available for: DS8000 and storage systems that run IBM Spectrum Virtualize.

Tier O Flash Capacity (GiB)

The total amount of storage space on the Tier 0 flash solid-state drives that can be used by Easy Tier for retiering the volume extents in the pool.

Available for: DS8000 and storage systems that run IBM Spectrum Virtualize.

Tier 1 Flash Available Capacity (GiB)

The amount of storage space that is available on the Tier 1 flash, read-intensive solid-state drives that can be used by Easy Tier for retiering the volume extents in the pool.

Available for: DS8000 and storage systems that run IBM Spectrum Virtualize.

Tier 1 Flash Capacity (GiB)

The total amount of storage space on the Tier 1 flash, read-intensive solid-state drives that can be used by Easy Tier for retiering the volume extents in the pool.

Available for: DS8000 and storage systems that run IBM Spectrum Virtualize.

Tier 2 Flash Available Capacity (GiB)

The available capacity on Tier 2 flash, high-capacity drives in the pool. Easy Tier can use these drives to retier the volume extents in the pool.

Available for: DS8000 storage systems.

Tier 2 Flash Capacity (GiB)

The total capacity on Tier 2 flash, high-capacity drives in the pool. Easy Tier can use these drives to retier the volume extents in the pool.

Available for: DS8000 storage systems.

Total Capacity Savings (%)

(Previously known as Total Data Reduction Savings) The estimated amount and percentage of capacity that is saved by using data deduplication, pool compression, thin provisioning, and drive compression, across all volumes in the pool.

The following formula is used to calculate the amount of storage space that is saved:

Provisioned Capacity - Used Capacity

The following formula is used to calculate the percentage of capacity that is saved:

((Provisioned Capacity - Used Capacity) \div Provisioned Capacity) \times 100

Available for: FlashSystem A9000 and FlashSystem A9000R, IBM Spectrum Accelerate, XIV storage systems with firmware version 11.6 or later, and resources that run IBM Spectrum Virtualize.

Unmapped Capacity (GiB)

(Previously known as Unassigned Volume Space) The total amount of space in the volumes that are not assigned to hosts. For Hitachi VSP non-thin provisioning pool space, this value is the sum of unassigned regular host-accessible volumes. Volumes that are used for thin-provisioning (pool volumes) are not included.

Available for: All storage systems.

Available Volume Capacity (GiB)

(Previously known as Effective Unallocated Volume Space) The total amount of remaining capacity that can be used by the existing volumes in the pools. The following formula is used to calculate this value:


Provisioned Capacity - Used Capacity

The capacity that is used by thin-provisioned volumes is typically less than their provisioned capacity. Therefore, the available capacity represents the difference between the provisioned capacity and the used capacity for all the volumes in the pool. For Hitachi VSP non-thin provisioned pool capacity, the unused volume capacity is always zero.

Available for: All storage systems, except FlashSystem A9000 and FlashSystem A9000R.

Used Capacity (%)

(Previously known as Physical Allocation)



The percentage of physical capacity that is used by the volumes in the pool, including the volumes in child pools. This value is always less than or equal to 100% because you cannot allocate more physical space than is available in a pool. Check the value for used capacity to see:

- Whether the physical capacity of the pool is fully allocated. That is, the value for used capacity is 100%.
- Whether you have sufficient capacity to provision new volumes with storage
- Whether you have sufficient capacity to allocate to the compressed and thin-provisioned volumes in the pool

Available for: All storage systems, except FlashSystem A9000 and FlashSystem A9000R.

Used Capacity (GiB)

(Previously known as Allocated Space) The amount of physical capacity that is used by the volumes in the pool. If the pool is a parent pool, the amount of space that is used by the volumes in the child pools is also included.

The capacity that is used by thin-provisioned volumes is less than their provisioned capacity, which is shown in the Provisioned Capacity column. If a pool does not contain thin-provisioned volumes, this value is the same as Provisioned Capacity.

Available for: All storage systems, except FlashSystem A9000 and FlashSystem A9000R.

Used Written Capacity (%)

(Previously known as Effective Used Capacity) For devices with inline hardware compression, the effective used capacity percentage is the percentage of capacity that is provisioned to the standard-provisioned volumes and the thin-provisioned volumes, given the drive compression savings.

Used Written Capacity (GiB)

(Previously known as Effective Used Capacity) For devices with inline hardware compression, the effective used capacity is the total amount of provisioned capacity that is used by all the volumes given the drive compression savings.

Written Capacity Limit (GiB)

(Previously known as Effective Capacity) For devices with inline hardware compression, the effective capacity is the amount of provisioned capacity that can be created in a storage system or pool given

the current data reduction savings. The effective capacity equals the usable capacity divided by the data reduction savings percentage.

Volume capacity metrics

You use the capacity chart to detect capacity shortages for the following types of volumes:

- · Space-efficient volumes such as compressed volumes and thin-provisioned volumes
- Standard-provisioned volumes that use Easy Tier to re-tier volume extents

You can review the capacity usage by space-efficient volumes to detect capacity shortfalls. You can also review the capacity usage of volumes that use Easy Tier to distribute volume extents across Enterprise HDD, Nearline HDD, and SSD drives.

To detect capacity shortages and investigate capacity usage trends, you can add the following metrics to the chart for volumes:

Capacity (GiB)

The capacity of the compressed or the thin-provisioned volume, which comprises the sum of the used and available capacity. For thin-provisioned volumes in XIV systems pools or IBM Spectrum Accelerate pools, capacity is the physical ("hard") capacity of the volume.

Available for: All storage systems.

Compression Savings (%)

The estimated amount and percentage of capacity that is saved by using data compression.

The following formula is used to calculate the amount of storage space that is saved:

```
written capacity - compressed size
```

The following formula is used to calculate the percentage of capacity that is saved:

```
((written capacity - compressed size) \div written capacity) \times 100
```

Available for: FlashSystem A9000 and FlashSystem A9000R, IBM Spectrum Accelerate, XIV storage systems with firmware version 11.6 or later, and resources that run IBM Spectrum Virtualize.

Exception: For compressed volumes that are also deduplicated, on storage systems that run IBM Spectrum Virtualize, this column is blank.

Enterprise HDD Capacity (GiB)

The amount of volume capacity that Easy Tier has placed on Enterprise hard disk drives.

Available for: DS8000 and storage systems that run IBM Spectrum Virtualize.

Nearline HDD Capacity (GiB)

The amount of volume capacity that Easy Tier has placed on Nearline hard disk drives.

Available for: DS8000 and storage systems that run IBM Spectrum Virtualize.

SCM Capacity (GiB)

The amount of volume capacity that Easy Tier has placed on Storage Class Memory (SCM) drives.

Available for: IBM Spectrum Virtualize systems, such as FlashSystem 9100, FlashSystem 7200, and Storwize family storage systems that are configured with block storage.

Tier O Flash Capacity (GiB)

The amount of volume capacity that Easy Tier has placed on Tier 0 flash drives.

Available for: DS8000 and storage systems that run IBM Spectrum Virtualize.

Tier 1 Flash Capacity (GiB)

The amount of volume capacity that Easy Tier has placed on Tier 1 flash, read-intensive drives.

Available for: DS8000 and storage systems that run IBM Spectrum Virtualize.

Tier 2 Flash Capacity (GiB)

The amount of volume capacity that Easy Tier has placed on Tier 2 flash, high-capacity drives.

Available for: DS8000 storage systems.

Used Capacity (GiB)

(Previously known as Allocated Space) The amount of space that is used by the compressed, thin-provisioned, or the Easy Tier volume. Typically, the space that is used by the compressed or thin-provisioned volume is less than the capacity of the volume. For Easy Tier volumes, used capacity is the capacity that is used by the volume's extents on the Enterprise HDD, Nearline HDD, or SSD drives.

Available for: All storage systems.

Written Capacity (GiB)

(Previously known as Written Space) The amount of data that is written from the assigned hosts to the volume before compression or data deduplication are used to reduce the size of the data. For example, the written capacity for a volume is 40 GiB. After compression, the volume used space, which reflects the size of compressed data that is written to disk, is just 10 GiB.

Key capacity concepts

Understand the terms that are used to measure and plan capacity.

Scenarios are used to explore the following capacity concepts for block storage systems:

- "Standard provisioning" on page 381
- "Thin-provisioning" on page 383
 - "Thin-provisioning savings" on page 383
 - "Data reduction savings" on page 383
 - "Total capacity savings" on page 384
 - "Effective capacity" on page 385


Tip: To review the capacity terms and definitions that are used, click "Definitions for the key capacity terms and concepts" on page 386.

Standard provisioning

You add Storage System A with a capacity of 100 TiB to IBM Spectrum Control for monitoring.

Then, you add a standard-provisioned volume with a capacity of 20 TiB.

You create standard-provisioned volumes when you want to dedicate the usable capacity of the storage system to the device that writes to it. From the perspective of the pool and storage system, the capacity is used and is no longer available to the storage system and pool.



The capacity values that are reported by the storage system for standard-provisioned volumes are as follows:

- **Used Capacity (%)** is reported as fully allocated, and **Used Capacity (GiB)** has the same value as the provisioned capacity of the volume.
- Available Capacity (%) is reported as 00.00.

The capacity of the data that is written to standard-provisioned volumes is not measured by the storage system. However, the written capacity of standard-provisioned volumes is included in the calculation of the total written capacity of the volumes. Because the data that is written to standard-provisioned volumes isn't reduced when data is stored, the written capacity has the same value as the provisioned capacity of the volume.



Figure 4. Overview chart: Written capacity of standard and thin-provisioned volumes

Tip: From the **Storage** menu, click **Block Storage Systems**. Double-click a storage system and check the **Provisioned Capacity** chart.


Thin-provisioning

Unlike standard-provisioned volumes, the capacity of thin-provisioned volumes is not dedicated to the device that writes to it. Devices get the capacity when they write data to the thin-provisioned volumes. Depending on the data reduction techniques that are supported by the storage system, the data that is written to the thin-provisioned volumes can be reduced before it is stored on the volumes.

Thin-provisioning savings

Thin-provisioning savings are the total amount of capacity that is saved in a pool, system, or volume by using capacity when needed as a result of write operations. The capacity that is saved is the difference between the provisioned capacity and the written capacity.

In this scenario, you add four thin-provisioned volumes with a total capacity of 200 TiB to Storage System B, which has a capacity of 100 TiB.



Thin-provisioning savings are the difference between the 200 TiB provisioned to the thin-provisioned volumes and the 150 TiB of data that is written to the volumes, which is 50 TiB.


(To enlarge, right-click the video and select the full screen or new tab option.)

View a short animation about measuring the capacity data reduction savings of thin-provisioned volumes.

Data reduction savings

Data reduction savings are the total amount of usable capacity that is saved in a pool, system, or volume through the application of data reduction algorithms on the written data, such as compression and deduplication. This saved capacity is the difference between the written capacity and the used capacity.

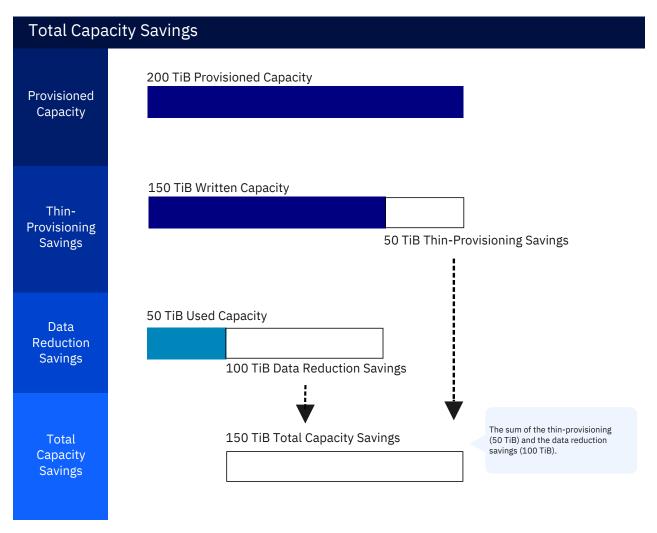
Before data is stored on thin-provisioned volumes, it can be reduced. For example, if your storage system supports compression and deduplication, these data reduction techniques are applied to reduce the amount of capacity that is needed to store the data.



In this scenario, the data reduction savings are the difference between the written capacity (150 TiB), and the used capacity (50 TiB), which is 100 TiB or 67%.

Total capacity savings

Total capacity savings are the amount of capacity that is saved in a pool, system, or volume through thin-provisioning and data reduction techniques.




In the scenario illustrated above, the thin-provisioning savings are (200 TiB - 150 TiB), which is 50 TiB. The data reductions savings are (150 TiB - 50 TiB), which is 100 TiB. So, the total capacity savings are 150 TiB.

Effective capacity

Effective capacity is the amount of provisioned capacity that can be created in a system or pool without running out of capacity given the current data reduction savings that are being achieved. To calculate effective capacity, you divide the capacity of the storage system or pool by the data reduction savings percentage.

Note: In some storage systems, restrictions in the system determine the maximum provisioned capacity that is allowed in a pool or system. In these cases, the effective capacity cannot exceed this limit.

You can use effective capacity to estimate how much capacity you can provision to your thin-provisioned volumes without running out of usable capacity.



In the scenario illustrated above, the storage system has a capacity of 100 TiB and a current data reduction savings of 66%. So, the effective capacity is 100/(1.00 - 0.67), which is 303 TiB. Currently, the provisioned capacity of the volumes is 200 TiB. Assuming that the same data reduction savings are achieved for subsequent write operations, more thin-provisioned volumes can be created with a total capacity of 303 TiB.

Definitions for the key capacity terms and concepts

These are capacity terms and concepts that are used:

Available capacity

The amount of usable capacity that is not yet used in a system, pool, array, or MDisk.

Capacity

The amount of usable capacity that is available for storing data on a system, pool, array, or MDisk after formatting and RAID techniques are applied.

Compression

A function that removes repetitive characters, spaces, strings of characters, or binary data from the data being processed and replaces characters with control characters. Compression reduces the amount of storage space that is required for data.

Data deduplication

A method of reducing storage needs by eliminating redundant data. Only one instance of the data is retained on storage media. Other instances of the same data are replaced with a pointer to the retained instance.

Data reduction

A set of techniques that can be used to reduce the amount of usable capacity that is required to store data. Examples of data reduction include data deduplication and compression.

Data reduction savings

The total amount of usable capacity that is saved in a pool, system, or volume through the application of data reduction algorithms on the written data, such as compression and deduplication. This saved capacity is the difference between the written capacity and the used capacity.

Effective capacity

The amount of provisioned capacity that can be created in a system or pool without running out of usable capacity given the current data reduction savings that are being achieved. This capacity equals the usable capacity that is divided by the data reduction savings percentage. In some storage systems, restrictions in the system determine the maximum provisioned capacity that is allowed in a pool or system. In those cases, the effective capacity cannot exceed this limit.

Provisioned capacity

The total capacity of all volumes in a pool or system.

Standard-provisioned volumes

Unlike thin-provisioned volumes, which use capacity when it is needed, the capacity that is provisioned to standard-provisioned or thick volumes is no longer available to the pool. Standard-provisioned volumes are fully allocated, and their capacity is reported as used because their capacity is no longer available to the pool and storage system.

Thin-provisioning

The ability to defer capacity allocation on a storage resource until data is actually written to it.

Thin-provisioning savings

The total amount of usable capacity that is saved in a pool, system, or volume by using usable capacity when needed as a result of write operations. The capacity that is saved is the difference between the provisioned capacity minus the written capacity.

Total capacity savings

The total amount of usable capacity that is saved in a pool, system, or volume through thin-provisioning and data reduction techniques. This saved capacity is the difference between the used usable capacity and the provisioned capacity.

Used capacity

The amount of capacity that is taken up by data in a system, pool, array, or MDisk after data reduction techniques have been applied.

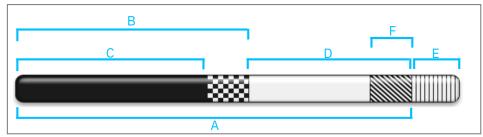
Written capacity

The amount of usable capacity that would have been used to store written data in a pool or system if data reduction was not applied.

Key storage values for pools

The metadata that is collected about storage systems is used to calculate key capacity values for pools. Use these values to monitor storage usage and to detect and identify capacity issues.

The following diagram shows the relationships between the key storage values for pools:



The following table describes the values on the diagram and how they are calculated:

Storage value	Formula	Description
Capacity		The total amount of physical and usable capacity in a pool.

Storage value	Formula	Description
Used Capacity	В	The amount of physical capacity that is used by the volumes in a pool, including the volumes in child pools.
Available Volume Capacity	PC - B	The amount of the Provisioned Capacity (PC) in the pool that is not used.
Used Capacity %	B ÷ A × 100	The percentage of physical capacity that is used by the volumes in a pool, including the volumes in child pools.
Provisioned Capacity %	PC ÷ A × 100	The percentage of the physical capacity that is committed or reserved for the provisioned capacity of the volumes in the pool.
Reserved Volume Capacity	B – C	The amount of pool capacity that is reserved but has not been used yet to store data on the thin-provisioned volume.
Provisioned Capacity (PC)	Sum of volume capacities	The total storage capacity on all the volumes in a pool, which includes thin-provisioned and standard-provisioned volumes.
Overprovisioned Capacity %	E = PC - A	The amount of capacity by which the Provisioned Capacity exceeds the physical capacity of a pool.
Shortfall %	((E ÷ (D + E)) × 100 E = Overprovisioned Capacity %	The percentage of the remaining capacity in a pool that is not available to be used. The higher the percentage, the more critical the shortfall of pool capacity.
Reserved Capacity	F	The amount of unused capacity in a pool that is reserved by provisioning and optimization tasks. Pool capacity is reserved when a provisioning or optimization task is created, and used when the task is run.

Capacity metrics for file storage systems

To review trends in capacity and space usage for file storage systems, you add metrics to capacity charts. You use the charts for filesets, file systems, and file system pools to detect capacity shortages and space usage trends.

Alphabetical lists of the capacity and space usage metrics that you can add to charts are provided in the following sections:

- "Capacity metrics for filesets" on page 388
- "Capacity metrics for file systems" on page 389
- "Capacity metrics for file system pools" on page 389

Capacity metrics for filesets

If sufficient data is collected about the filesets in your file storage systems, you can view charts that compare the used capacity of the filesets. The following metric is displayed in the capacity chart for filesets:

Used Capacity (GiB)

The amount of storage capacity that is used by the fileset. Used capacity is not provided for filesets that are cache targets.

Available for: All storage systems.

Capacity metrics for file systems

To detect capacity shortages and investigate capacity usage trends, you can add the following metrics to the capacity chart for file systems:

Available Inodes

The number of inodes that are available in a file system.

Available Capacity (GiB)

The amount of storage capacity that is available (not allocated) on a file system.

Capacity (%)

The percentage of the total storage space on the file system that is used by files and directories.

Available for: All storage systems.

Capacity (GiB)

(Previously known as Total Capacity) The capacity on the file system.

Available for: All storage systems.

Maximum Inodes

The total number of inodes that the file system can contain. This value consists of the available inodes and the used inodes for the file system.

Available for: All storage systems.

Physical Capacity (GiB)

The raw capacity of the partition where a file system resides.

Used Inodes

The number of used inodes on a file system. An inode is the internal structure that describes the individual files or directories in the file system metadata. An inode contains the node, type, owner, and location of a file or directory.

Available for: All storage systems.

Used Inodes (%)

The percentage of inodes that are already used on the file system. Each time that you create a file or directory on the file system, an inode is allocated to the file or directory.

Available for: All storage systems.

Internal Used Capacity (GiB)

The amount of storage space that is unavailable (allocated) on a file system. For IBM Spectrum Scale, this value does not include the space that is used by migrated data on external pools.

Capacity metrics for file system pools

If sufficient data is collected about the file system pools in your data center, you can view charts that compare the capacity of the pools with the space that is allocated to the pools and the space that is still available in the pools.

The internal resource for IBM Spectrum Scale storage systems is called **Pools**. For all other file storage systems, the internal resource is called **File System Pools**.

To detect capacity shortages and investigate space usage trends, you can add the following metrics to the capacity chart for file system pools:

Available Capacity (GiB)

The amount of unallocated storage capacity in the pool. Available capacity usually comprises the space that can be used for storage. However, if the pool is not formatted, the amount of overhead capacity might also be included in the calculation.

Available for: All storage systems.

Capacity (%)

The percentage of space in the file system pool that is being used.

Available for: All storage systems.

Capacity (GiB)

(Previously known as Total Capacity) The capacity in the storage pool.

Available for: All storage systems.

Used Capacity (GiB)

The amount of space in the file system pool that is being used.

Available for: All storage systems.

Capacity metrics for object storage systems

To review trends in capacity and space usage for object storage systems, you add metrics to capacity charts. Use the charts for containers to detect capacity shortages and space usage trends for the containers in your object storage systems.

Capacity metrics for containers

If sufficient data is collected about the containers in your object storage systems, you can view charts that compare the used capacity of the containers. You can add the following metrics to the capacity chart for containers:

Available Objects

The number of objects that you can add to the container on the object storage system.

Available for: IBM Spectrum Scale storage systems.

Available Capacity (GiB)

The amount of file system capacity that can be used to store object data for the container.

Available for: IBM Spectrum Scale storage systems.

Capacity Quota (%)

The percentage of the capacity quota for the container that was used when the storage system was last probed. The quota limits the amount of file system capacity that can be used by the container. This column contains a value only if a capacity quota is defined for the container on the OpenStack Swift object storage system.

If the used capacity value for the container exceeds the capacity quota value, the percentage exceeds 100%. The used capacity can exceed the capacity quota if the quota is set after object files are uploaded to the file system or if an object file without size information is uploaded to the file system.

Available for: IBM Spectrum Scale storage systems.

Capacity Quota (GiB)

The amount of file system capacity that can be used to store object data for the container. This column contains a value only if a capacity quota is defined for the container on the OpenStack Swift object storage system.

Available for: IBM Spectrum Scale storage systems.

Objects

The number of objects in the container on the object storage system.

An object stores data content, such as files, videos, images, virtual machine snapshots, and other unstructured data. In IBM Spectrum Scale, objects are stored as files on the GPFS file system.

Available for: IBM Spectrum Scale storage systems.

Objects Quota

The number of objects that can be stored in the container. This column contains a value only if an objects quota is defined for the container on the OpenStack Swift object storage system.

Available for: IBM Spectrum Scale storage systems.

Objects Quota (%)

The percentage of the objects quota for the container that was used when the storage system was last probed. The quota limits the number of objects that can be stored in the container. This column contains a value only if an objects quota is defined for the container on the OpenStack Swift object storage system.

If the number of objects that are used exceeds the objects quota, the percentage exceeds 100%. The number can exceed the quota if the quota is configured after the object files are stored on the file system.

Available for: IBM Spectrum Scale storage systems.

Used Capacity (GiB)

The amount of file system capacity that is used by the objects in the container. If a probe is run immediately after objects are added to a container, the used capacity value might not reflect the updates until after the next scheduled probe is run.

Available for: IBM Spectrum Scale storage systems.

Capacity metrics for tiers

Review trends in the capacity and space usage for tiers. You use the charts to detect capacity shortages and space usage for the tiers in your storage environment.

Based on historical space usage, you can plan the capacity growth for tiers.

You can view values for the following metrics in the tier charts:

Capacity (GiB)

The total amount of storage space in the tier, which comprises the total storage space of the pools in the tier.

Used Capacity (GiB)

The capacity that is used by pools in the tier.

Monitoring and administering applications, departments, and general groups

You create applications to monitor the storage capacity, space usage, and performance of applications, and you create departments to monitor the space usage of the applications in the department. You can structure your applications and departments hierarchically to match the structure of your business organization.

You can map the application to the storage resources that the application uses to do capacity trending, health monitoring, and performance troubleshooting tasks. For example, a department might use 15 applications and be part of another five departments. Also a department might share storage resources with another department, subdepartment or application even if they do not belong in the same business hierarchy. Another example, might be if you are contacted because an application within a department is experiencing a performance issue. You carry out tasks like capacity reporting or trending on the behalf of the application.

Create general groups to quickly view information about storage resources that have common characteristics. For example, you might group the subset of ports on a SAN Volume Controller that are used for inter-node communication, or all the storage systems with lease agreements that end in the current year, or the storage systems that are used by a critical business application.

Application administration

Go to the pages for creating and administering the applications that you want to monitor.

Table 212. Create and administer applications and subcomponents	
Actions	Navigation
Create applications and assign storage resources directly to the application.	 In the menu bar, go to Groups > Applications. On the Applications page, click Create Application. On the Select a Method page, click Assign storage resources to the application.
Create subcomponents that can have its own storage resources.	 In the menu bar, go to Groups > Applications. On the Applications page, click Create Application. On the Select a Method page, click Create a subcomponent that can have its own set of storage resources.
Create a filter to use in assigning resources.	 In the menu bar, go to Groups > Applications. On the Applications page, right-click a row in the Applications list and click View Details. In the Members section, click Filters. On the Filters page, click Create Filter.
Add applications as subcomponents to an application.	 In the menu bar, go to Groups > Applications. Right-click the applications that you want to add as subcomponents, and then click Add to > Application. On the Add to Application page, select the applications and click Save.
Add applications to departments.	 In the menu bar, go to Groups > Applications. On the Applications page, select one or more applications, right-click and then click Add to > Department. On the Add to Department page, select the applications that you want to add as members and click Save.
Add applications to subdepartments.	 In the menu bar, go to Groups > Departments. Right-click the department, and then click View Details. In the General section, click Subdepartments. Right-click the subdepartment, and click View Details. In the General section, click Applications. Click Create Application.

Table 212. Create and administer applications and subcomponents (continued)		
Actions	Navigation	
Remove applications and subcomponents that you do not want to monitor.	 In the menu bar, go to Groups > Applications. On the Applications page, select one or more applications, and click Delete. If you select Delete subcomponents, the subcomponents that belong only to that application are removed. 	
Remove subcomponents within direct context of an application that you do not want to monitor.	 In the menu bar, go to Groups > Applications. Right-click the application and click View Details. In the General section, click Subcomponents. Right-click the subcomponent, and click Remove from Application. 	
Remove selected resources from applications or subcomponents	 In the menu bar, go to Groups > Applications. Right-click the application and click View Details. In the Related Resources section, click an eligible resource. Optional: In the Members section, click the storage resource that is part of the application. On the resource page, right-click the resource, and click Remove from Application. 	
Troubleshoot applications	In the menu bar, go to Groups > Applications . Right-click the application, and then click View Details . Investigate the performance of the application and the storage resources that are associated with the application by viewing the charts on the Overview page. You can investigate the performance of related resources by clicking the resource and clicking the Performance tab.	
Troubleshoot subcomponents	 In the menu bar, go to Groups > Applications . Right-click the application, and then click View Details. In the General section, click Subcomponents. Right-click the subcomponent, and click View Details. Investigate the performance of the storage resources that are associated with the subcomponent by viewing the charts on the Overview page. You can investigate the performance of related resources by clicking the resource and clicking the Performance tab. 	

Department administrationGo to the pages for creating and administering the departments that you want to monitor.

Table 213. Create and administer departments and subdepartments	
Actions	Navigation
Create departments and associate members (applications and departments) to create a hierarchal structure.	 In the menu bar, go to Groups > Departments. On the Departments page, click Create Department. On the Add Members page, click Add one or more applications or Create subdepartment.
Add departments as subdepartments to a department.	In the menu bar, go to Groups > Departments . On the Departments page, right-click the departments that you want to add as subdepartments, and then click Add to Department .
Add applications to departments.	In the menu bar, go to Groups > Departments. Right-click the department, and then click View Details. In the General section, click Applications. Click Create Application.
Add applications to subdepartments.	 In the menu bar, go to Groups > Departments. Right-click the department, and then click View Details. In the General section, click Subdepartments. Right-click the subdepartment, and click View Details. In the General section, click Applications. Click Create Application.
Remove departments, subdepartments, and applications, that are members of the department.	 In the menu bar, go to Groups > Departments. On the Departments page, right-click one or more departments and click Delete. If you select Delete subdepartments and applications, the subdepartments and applications that belong only to the department, that is being removed, are removed.

Table 213. Create and administer departments and subdepartments (continued)	
Actions	Navigation
Remove applications from a department or subdepartment.	 In the menu bar, go to Groups > Departments. Right-click the department, and then click View Details. In the General section, click Subdepartments. Right-click the subdepartment, and click View Details. In the General section, click Applications. Select and right-click the applications, and then click Delete or Remove from Department.
Remove subdepartments within direct context of a department that you do not want to monitor.	 In the menu bar, go to Groups > Departments. Right-click the department, and then click View Details. In the General section, click Subdepartments. Right-click the subdepartments that you want to remove, and then click Remove from Department.

General group administrationAdminister the general groups that you use to organize storage resources with common characteristics.

Table 214. Administering general groups	
Actions	Navigation
Create a general group and add resources.	In the menu bar, click the resource list page for the resources that you want to add. For example, if you want to add file storage systems, click Storage > File Storage Systems.
	Right-click one or more resources and click Add to General Group .
	3. Click Add to new group , specify a name and description for the new group, and click Save .
Add resources to general groups.	In the menu bar, click the resource list page for the resources that you want to add. For example, if you want to add block storage systems, click Storage > Block Storage Systems.
	Right-click one or more resources and click Add to General Group .
	3. Take one of the following actions:
	 a. To add the resources to one or more existing groups, click Add to existing groups, click the appropriate groups in the list, and click Save.
	 b. To add the resources to a new group, click Add to new group, specify a name and description for the new group, and click Save.

Table 214. Administering general groups (continued)		
Actions	Navigation	
Set which alert policy manages a general group <u>Learn more</u>	 Go to the details page for the general group for which you want to set the policy. In the menu bar, click Groups > General Groups. Right-click one of the general groups, then click View Details. Click Alerts in the General section. Click Set Policy from the Policy Actions menu. To set the alert policy for multiple general groups, go to Settings > Alert Policies. Double-click the policy, click the resources tab, then click Edit Resources. 	
Create an alert policy from the alert definitions and notification settings in a general group	 Go to the details page for the general group from which you want to create the policy. In the menu bar, click Groups > General Groups. Right-click one of the general groups, then click View Details. Click Alerts Definitions in the General section. Click Create Policy from the Policy Actions menu. 	
View or modify the alert policy that manages a general group Learn more about defining alerts and notification settings	 Go to the details page for the general group whose policy you want to view. In the menu bar, click Groups > General Groups. Right-click one of the general groups, then click View Details. Click Alerts in the General section. Click View Policy from the Policy Actions menu. 	
Delete general groups.	 In the menu bar, click Groups > General Groups. Right-click one or more general groups and click Delete. To delete subgroups, ensure that the Delete subgroups? check box is selected. When a group is deleted, the information that is associated with the group, such as alerts and alert definitions, is removed from the product. The resources that were members of the group are removed from the group but are still monitored by the product. 	
Remove resources from a general group.	 In the menu bar, click Groups > General Groups. Right-click a general group and click View Details. In the Members section, click the type of resource that you want to remove. For example, if you want to remove some of the volumes from the group, click Volumes. Right-click the resources that you want to remove and click Remove from General Group. Click OK to confirm that you want to remove the resources from the general group and from all of its subgroups. 	

Table 214. Administering general groups (continued)	
Actions	Navigation
Add general groups as subgroups.	 In the menu bar, click Groups > General Groups. Right-click one or more general groups and click Add to General
	Group. 3. Take one of the following actions:
	a. To add the groups as subgroups of existing groups, click Add to existing groups , click one or more groups in the list, and click Save .
	 b. To add the groups as subgroups of a new group, click Add to new group, specify a name and description for the new group, and click Save.
Remove subgroups from a group	1. In the menu bar, click Groups > General Groups .
hierarchy.	2. Right-click a general group and click View Details .
	3. Click Subgroups in the General section.
	4. Right-click the general groups that you want to remove and click Remove from General Group .
	When you remove a subgroup from its parent group, the subgroup is moved to the same level in the hierarchy as the parent group. The subgroup is still monitored by the product.
Create a subgroup.	1. In the menu bar, click Groups > General Groups .
	2. Right-click the general group that you want to create a subgroup for and click View Details .
	3. Click Subgroups in the General section.
	4. Click Create Subgroup .
	5. Specify a name and description for the new group.
	6. Customize the icon for the group by clicking the existing icon and selecting another icon.
	7. Click Create .
Delete subgroups.	1. In the menu bar, click Groups > General Groups .
	Right-click the general group that contains the subgroups that you want to delete and click View Details .
	3. Click Subgroups in the General section.
	4. Right-click one or more subgroups and click Delete .
	5. To delete further subgroups, ensure that the Delete subgroups? check box is selected.
	When a group is deleted, the information that is associated with the group, such as alerts and alert definitions, is removed from the product. The resources that were members of the group are removed from the group but are still monitored by the product.

Related tasks

Monitoring capacity usage at different levels of a business hierarchy

The monitoring and management of applications and departments enables you to monitor storage capacity usage, recognize trends, monitor health status, and troubleshoot performance of the storage resources in your business organization.

Comparing storage usage in each department

In this tutorial, you compare the storage usage in each department within your organization to spot a potential storage usage issue.

Using applications and subcomponents to monitor capacity and space usage

To monitor the performance, capacity, and space usage of the applications in your business organization, create applications and subcomponents. You can also add applications to departments so that storage capacity and usage can be monitored in an overall business hierarchal manner.

Viewing storage capacity and usage trends

You can view the capacity and usage trends of storage resources used by an application and subcomponents to diagnose how much storage a particular application is using in relation to the other applications and departments in the business hierarchy.

Applications

View information about storage area network resources that you can model by using applications in your business environment. The monitoring and management of applications enables you to perform capacity trending, health monitoring, and performance troubleshooting tasks.

An application can be part of a department and have subcomponents that can contain a five level deep hierarchy. Applications that are grouped together can range from large line of business systems to specialized software, in a department that runs on either client computers or servers. For example, an application might be an automated billing system in the Finance department, VMware running in the Information Technology department or an email marketing system that is part of the Marketing department.

Another example might be an application named *Database* and the database contains logs and data, and a best practice is to place the logs and data on separate volumes. You might create subcomponents for the logs and data and these subcomponents become part of the *Database* application.

A storage administrator can use applications for planning purposes to diagnose how much storage a particular application is consuming in relation to the other applications in the business hierarchy. Another applicable usage for applications might be troubleshooting performance issues. The administrator checks the related resources that make up the application to see if there is a health status issue or can examine the performance of volumes that are used by the application.

The **Applications** page shows capacity information about the applications that are monitored and the actions that you can use to view and manage them. To see more detailed information, view properties and details for the application. If the application is associated with a department, information about that department is also shown.

The following information shows the resources and objects that can be added to the application.

Table 215. Resources and objects that can be added to the application	
Object	Resources or objects that can be added to the application
Application	 Applications Data stores Hypervisors File Systems (Only file systems that are monitored through a Storage Resource Agent.) Filesets Servers Shares Vaults (IBM Cloud Object Storage) Virtual Machines (Only can be added with the CLI modifyappgroup command) Volumes Volume Groups

Related tasks

Defining alerts for applications

Define alerts for changes in the configuration, attributes, and performance of the servers, volumes, filesets, and shares in your application.

Creating an application hierarchy

To monitor, the capacity, space usage, and performance of the storage resources that your applications use, create application models.

In this scenario, you want to monitor the capacity, space usage, and performance of a DB2 application.

The DB2 application has two subcomponents, which you want to monitor separately:

- logs
- data

To monitor the DB2 applications and the data and logs subcomponents, complete these tasks:

- Create an application that is named DB2.
- Create applications that are named data and logs
- Associate storage resources with the data and logs applications. For example, in this case, associate the volumes that the data application uses, type data_vol*. For the logs application, type logs_vol* in the Volumes field of A selection of volumes that match a pattern filter. Then click Belonging to and click Storage System. Type the name of the storage system or the pattern that matches the name of the storage system.
- Add the data application and the logs application as subcomponents to the DB2 application.

Tip:

When you create complex application models for applications, it is not necessary to create a filter to associate storage resources with the application because the application inherits the storage resources that are associated with its subcomponents. However, if you want to monitor storage resources that are not associated with the subcomponents, you must create a filter that associates the additional storage resources that you want to monitor with the application.

You can now view information about the capacity, space usage, and performance of the storage resources on the Overview page for the application and subcomponents.

You can create departments and add applications to the departments.

Related concepts

"Application and department hierarchies" on page 411

The interactive breadcrumb on the Overview page provides a visual representation of the hierarchical relationships for applications and departments and enables you to navigate quickly between the elements in the hierarchy.

Creating applications

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies for applications, you can create applications for monitoring in IBM Spectrum Control.

Use the GUI to create applications, assign resources, and subcomponents for monitoring by IBM Spectrum Control.

To determine the storage capacity usage by the storage resources that the applications use, you create a data model of the applications for your storage environment.

To create an application, complete these steps:

- 1. In the menu bar, go to **Groups** > **Applications**.
- 2. On the Applications page, click Create Application.
- 3. On the **Create Application** page, enter information about the application.

Note: You cannot have a duplicate application name on the same level in your application hierarchy. For example, you create an application named *Finance* and create a subcomponent named *Business*. You can create another application named *Business* as a top level application, but you cannot add the new *Business* application as a subcomponent of the *Finance* application.

You can customize the icon for your application by clicking the existing icon on the **Create Application** page and selecting another icon from the **Customize Icon** page.

- 4. Click Create.
- 5. Optional: On the **Select a Method** page, select **Assign storage resources to the application that was created** and follow the instructions to add the resources directly to the application.
- 6. Click **Preview** to see a list of filter matches before you finalize your filtering criteria, otherwise, click **Save**.
- 7. Optional: Select **Create a subcomponent that can have its own set of storage resources** and follow the instructions to create a subcomponent of the application.
- 8. Click Create.

When the application creation is created, you can view status information, block and file capacity and additional data on the **Applications** page and in properties.

If the View Performance action is available, you can view the aggregated information for all of the volumes in the application on the performance view. By default, the Total I/O Rate and the Total Data Rate are displayed in the performance chart. You can also select to display the read and write I/O rates, and the read and write data rates in the chart.

Note: The application performance feature isn't available to all IBM Spectrum Control users. Look for news from IBM about when this feature will be available to you.

For a detailed view of the application and related resources, on the **Applications** page, right-click the application and select **View Details**. To see the subcomponents for the application, click **Subcomponent** on the application details page.

Related tasks

Adding subcomponents

You created a new application and a subcomponent, with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

Using filters to assign resources to applications

Instead of manually adding the resources that you want to assign to new or existing applications, you can create filters to assign resources based on name pattern matching.

Viewing information about applications

You can view information about the application that you created as related resources, filters that are used to populate the application with resources, any subcomponents, capacity trending and performance information.

Using filters to assign resources to applications

Instead of manually adding the resources that you want to assign to new or existing applications, you can create filters to assign resources based on name pattern matching.

When you create a filter to assign storage resources to applications, the resources that match the criteria that are defined in the filter are automatically added to the application or subcomponent (member application).

The use of wildcard characters enables you to customize and expand your search to add related resources. The asterisk (*) and the question mark (?) can be used as wildcard characters. The asterisk matches one or more characters. The question mark matches a single character. Use a comma-separated list to add multiple resources. As new resources are added or modified and they match the existing filter, they automatically become part of the application. For example, a new server is added, a volume is renamed or a share is mounted on a server.

If you edit a filter, the resources that no longer match the filter are removed from the application and the resources that now match the filter are added to the application.

If you remove the filter, all storage resources that match the filter criteria are removed from the application.

Filtering is extremely useful in storage environments where resources are predefined and named in a consistent format. Instead of adding each resource separately, you create filtering criteria that matches your naming format in your environment. You can use filtering when you create a new application or directly from the details page for an existing application.

You can only select one resource type filter at a time during the creation workflow, but you can create multiple filters for applications.

Tip: For any volume that you add to an application, the performance information for that volume, including historical information, is aggregated with the information for the existing volumes.

Note: The application performance feature isn't available to all IBM Spectrum Control users. Look for news from IBM about when this feature will be available to you.

The filters are divided into the following categories:

All storage that belongs to a set of servers/hypervisors

You can add servers or hypervisors to the application or subcomponent (member application).

For example, in the **Define Name Patterns** section, enter *blade, in the **Server/Hypervisor** field and click **Save**. The filter adds all known servers or hypervisors that have a name that ends with blade to the application.

If you enter server???.abc.com, in the **Server/Hypervisor** field, the filter adds all servers, in the domain with names that begin with server followed by three characters and ends with .abc.com to the application. If you enter, accounting*.*, the filter adds all servers with names that begin with accounting regardless of the domain the system belongs to.

If you enter cs-blade47.lab.abc.com, cs-blade49.lab.abc.com or cs-blade47*, cs-blade49*, the filter adds multiple servers to the application.

A selection of volumes that match a pattern

You can add volumes of a certain name format, volumes that belong to a server or hypervisor, volumes from a specific storage system or volumes from a pool to the application or subcomponent. You can only select one resource option per filter creation.

For example, in the **Define Name Patterns** section, if you enter *vol, in the **Volumes** field, mark the **Belonging to** box, click **Server/Hypervisor**, enter *.lab.abc.com and click **Save**. The filter adds all volumes that have a name that ends with vol and are mapped to servers that have a name that ends with .lab.abc.com to the application.

If you enter *vol, in the **Volumes** field, mark the **Belonging to** box, click **Storage System**, and enter ds8000. The filter adds all volumes that have a name that ends with vol that belong to the ds8000 storage system.

A selection of shares that match a pattern

You can add shares and exports that are mounted on a server or hypervisor, or shares and exports that are exported from a specific storage system to the application or subcomponent. You can select only one resource option per filter creation.

For example, in the **Define Name Patterns** section, if you enter *share in the **Shares** field, mark the **Belonging to** box, click **Server/Hypervisor**, enter *.lab.abc.com and click **Save**. The filter adds shares that have a name that ends with share that are mounted on servers that have a name that ends with .lab.abc.com.

If you enter share3 in the **Shares** field, mark the **Belonging to** box, click **Storage System** and enter netapp, the filter adds all the shares that are named share3 and are exported from the netapp storage system.

A selection of filesets that match a pattern

You can add filesets that are associated with file systems, from a storage system to the application or subcomponent. You can select only one resource option per filter creation.

For example, in the **Define Name Patterns** section, if you enter * in the **Filesets** field, mark the **Belonging to** box, click **Storage System**, enter scale* and click **Save**, the filter adds all the filesets that are exposed by storage systems that have names that begin with scale to the application.

If you enter * in the **Filesets** field, mark the **Belonging to** box, click **File System** and enter fs1, the filter adds all the filesets that are associated with the file systems named fs1 to the application.

A selection of vaults that match a pattern

You can add vaults of a certain name format and vaults that belong to a specific IBM Cloud Object Storage system.

For example, in the **Define Name Patterns** section, if you enter $*v_cloud*$ in the **Vault** field, select **Belonging to**, and enter *.lab.abc.com in the **Cloud Object Storage System** field, the filter adds all vaults that have a name that includes v_cloud and belong to IBM Cloud Object Storage systems with names that end in .lab.abc.com.

If you enter $\star v_cloud_ibm$ in the **Vaults** field, select **Belonging to**, and enter COS_storage in the **Cloud Object Storage System** field, the filter adds all vaults that have a name that ends with v_cloud_ibm and belong to the COS_storage storage system.

To create a filter, complete these steps:

- 1. In the menu bar, go to **Groups>Applications**.
- 2. On the **Applications** page, right-click a row in the Applications list and click **View Details**.
- 3. In the **Members** section, click **Filters**.
- 4. On the Filters page, click Create Filter.
- 5. On the **Create Filter** page, select a resource type filter and enter the necessary information.
- 6. Click **Preview** to see a list of filter matches before you finalize your filtering criteria.
- 7. On the **Filter Matches** page, click **Save** to create the filter.

The **Members** section of the details page shows the number of filters that you created, for example, **Filters (6)**. Click **Filters** to access the **Filters** page and see the filters that were created. For example, The number of matches: 26, which indicates there are 26 matches for the filter criteria. Click the link to see the resources for the filter matches that are available.

Tip: To view a description of a field on the **Create Filter** page, hover the mouse pointer over the help icon for that field.

Related tasks

"Creating applications" on page 400

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies for applications, you can create applications for monitoring in IBM Spectrum Control.

"Adding subcomponents" on page 403

You created a new application and a subcomponent, with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

Adding subcomponents

You created a new application and a subcomponent, with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

To create a subcomponent for an existing application, complete these steps:

- 1. In the menu bar, go to **Groups** > **Applications**.
- 2. On the Applications page, select one or more applications, right-click and then, click **Add to..** > **Application**.
- 3. On the **Add to Application** page, choose the applications that you want to become subcomponents of your selected applications.

The **Hierarchy** column shows the path of only the ancestor applications; it does not show departments.

- 4. Click Save.
- 5. Optional: On the details page for the application, click **Subcomponents**.
- 6. On the Subcomponents page, click Create Subcomponent.
- 7. On the **Create Subcomponent** page, enter information including name, description, type, and subtype.
- 8. Click Create.

After the subcomponent is created, you can view the results on the **Subcomponent** page and in the properties. For a more detailed view of the subcomponent, select the subcomponent from the **Subcomponent** page and click **Actions** > **View Details**.

Adding resources manually to applications

You can add resources directly from a resource page to your applications.

Only the following resources can be added directly to applications:

- · Data stores
- Hypervisors
- File Systems (Only server file systems that are monitored through a Storage Resource Agent.)
- Filesets
- Servers
- Shares
- Vaults (Cannot be added with the CLI)
- Virtual Machines (Only can be added with the CLI **modifyappgroup** command)
- Volumes
- Volume Groups

Tip: For any volume that you add to an application, the performance information for that volume, including historical information, is aggregated with the information for the existing volumes.

Note: The application performance feature isn't available to all IBM Spectrum Control users. Look for news from IBM about when this feature will be available to you.

To add a resource directly to an application, complete these steps:

- On a resource page, right-click a resource and click Add to Application.
 For example, to add a server to an application, in the menu bar, go to Groups > Servers. On the Servers page, select a server from the Servers list, right-click and click Add to Application.
- 2. On the **Add to Application** page, select one or more applications to assign the resource to.

3. Click Save.

The resource is added to the application and is displayed on the **Applications** page.

Related tasks

"Using filters to assign resources to applications" on page 401

Instead of manually adding the resources that you want to assign to new or existing applications, you can create filters to assign resources based on name pattern matching.

Adding resources with the command line interface

You can use the **modifyappgroupviafile** command-line interface (CLI) command to automate your bulk assignment of eligible storage resources to applications. To add a single resource to an application, use the **modifyappgroup** command.

Related information

modifyappgroupviafile command modifyappgroup command

Viewing information about applications

You can view information about the application that you created as related resources, filters that are used to populate the application with resources, any subcomponents, capacity trending and performance information.

To view the details of an application, use these steps:

- 1. In the menu bar, go to **Groups** > **Applications**.
- 2. On the **Applications** page, right-click an application and select **View Details**.

The details page shows information about the specific application. The information is available from the following sections on the page:

- **Overview**: The overall capacity trending about the application and performance information. An interactive breadcrumb that shows the hierarchy of applications and departments and a list of all the top level applications.
- **Properties**: The key information about the application.
- **Subcomponents**: The number of subcomponents that are a members of the selected application. If there is more than one subcomponent, the number of subcomponents is displayed.
- **Members**: The information on the resource types that were added directly to the application or subcomponent. The number next to the resource type, for example, Servers (2), shows the number of resources of that type in the application.
- **Filters**: The number of filters created and the filtering criteria used for the resources that were added to the application.
- **Related Resources**: The information about the resources that are related to the application.

Tip: To view descriptions of the information that is available on an application details page, click the **Help** icon in the upper-right corner of the page.

Related concepts

Application and department hierarchies

The interactive breadcrumb on the Overview page provides a visual representation of the hierarchical relationships for applications and departments and enables you to navigate quickly between the elements in the hierarchy.

Related tasks

Creating an application hierarchy

To monitor, the capacity, space usage, and performance of the storage resources that your applications use, create application models.

Removing applications and subcomponents

You can remove applications and their subcomponents. The data that was collected to monitor the applications is also removed from the database repository.

To remove applications and subcomponents from a level of the business hierarchy, complete the following steps:

- 1. In the menu bar, go to **Groups** > **Applications**.
- 2. On the **Applications** page, select one or more applications, and click **Delete**.
 - a) If you select **Delete Subcomponents**, the subcomponents that belong only to the application are removed.
 - b) If you do not select **Delete Subcomponents**, any subcomponents of the application, that is being removed, then become members of the parent application or department. If the application is a top level application, and, the subcomponents are not members of any other application, the subcomponents become top level applications.

The applications and subcomponents are removed and are no longer displayed on the **Applications** page.

To remove subcomponents from within the direct context of the application, use these steps:

- 1. In the menu bar, go to **Groups** > **Applications**.
- 2. On the Applications page, select an application from the Application list and right-click **View Details**.
- 3. Optional: Select a application from the **Applications** page and right-click **View Properties**.
 - a. Click the **Subcomponents** tab and continue with steps 5-6.
- 4. On the application details page, click **Subcomponents** in the General section.
- 5. On the Subcomponents page, select the subcomponent and right-click **Remove from Application**.
- 6. Confirm that you want to remove the subcomponent and click **OK**.

Removing resources from applications

Related resources within the context of an application can be removed from applications. The action is only available when viewing the related resources that belong to an application on the details page.

The following resources can be removed from an application:

- Volumes
- Shares
- Servers
- File systems (Only server file systems that are monitored through a Storage Resource Agent.)
- Filesets
- · Volume Groups
- Hypervisors
- Data Stores

Resources that were added through a filter match are removed if the filter is modified and no longer includes the resource, the filter is removed from IBM Spectrum Control, or the resource is removed from IBM Spectrum Control. Resources can be added back to the application if the filter is modified and includes the resource that was originally removed, a new filter is created that matches the resource, the resources are added to IBM Spectrum Control, or you use the **Add to Application** action.

For example, a server named, abc.xyz.com is added to an application named, ApplicationABC because it matched the server name filter, abc*. You decide to remove the abc.xyz.com server from the ApplicationABC application and remove the abc.xyz.com server from IBM Spectrum Control. Later you want to add the abc.xyz.com server back to IBM Spectrum Control. The abc.xyz.com server is added to the ApplicationABC application because the name matches the server resource filter that

you originally created for the ApplicationABC application. If you do not want the resource added back to the application, you can further refine your filters. Click the **View Filters** link on the confirmation dialog to make any adjustments to the filtering criteria before creation.

Resources that are automatically added to an application because they are related to resources that were explicitly added to the application *are not* removed from the application. The resources are automatically removed when the corresponding related resource that was added to the application, is also removed from the application.

For example, you add a server to the application that has one volume from a storage system. The volume is added as a related resource of that server in the application. If you attempt to remove the volume from the application, the volume is *not* removed since it is a related resource. The volume is only removed when the server, that was explicitly added to the application, is removed.

To remove resources directly from an application, use these steps

- 1. In the menu bar, go to **Groups > Applications**.
- 2. On the Applications page, select an application from the Application list and right-click View Details.
- 3. On the application details page, click an eligible resource in the Related Resource section.
- 4. On the resource page, select the resource from the list, and right-click **Remove from Application**.
- 5. Confirm that you want to remove the resource.
- 6. Click Remove.

The resources that were added to the application are removed from the application.

Tip: When you remove a member volume from an application, the performance information for that volume, including historical information, is no longer aggregated with the information for the existing volumes.

Note: The **View Performance** feature for applications isn't available to all IBM Spectrum Control users. Look for news from IBM about when this feature will be available to you.

Related tasks

"Using filters to assign resources to applications" on page 401

Instead of manually adding the resources that you want to assign to new or existing applications, you can create filters to assign resources based on name pattern matching.

Departments

View information about departments and structure the hierarchy to mirror your business organization. Monitor the detailed information about capacity and space usage that is collected for departments and subdepartments, and monitor the performance of the applications and the application subcomponents that are added to departments and subdepartments.

A department can be hierarchical in its organizational layout. For example, a department might use 15 applications and be part of another five departments. A department might share storage resources with another department, subdepartment or application even if they do not belong in the same hierarchy. For example, in a collaboration scenario, a single IBM SAN Volume Controller might be shared by multiple departments. You can provide reports on the storage capacity that is used at different levels of your business hierarchy to regulate and compare the capacity storage, and in turn, track the trending of storage usage by specific departments.

The **Departments** page shows storage capacity information about the top level departments, the subdepartments and any applications that belong to the department that are monitored by IBM Spectrum Control and the actions that you can use to view and manage the resources that are defined to them. To see more detailed information, view properties and details for the department.

The following information shows the objects that you can add using the Add to Department function.

Table 216. Objects that can be added to the department	
Object	Objects that can be added to the department
	ApplicationsDepartments
Department	

Creating a department hierarchy

To monitor the capacity and space usage of the department, you want to create a department hierarchy and associate applications with the department hierarchy. You can also monitor the performance of the storage resources that are associated with the applications that you add to the department hierarchy.

In this scenario, you want to monitor the capacity and space usage of departments in a business organization that sells books. You also want to monitor the performance of the storage resources that are associated with the applications or application subcomponents that are used by the departments.

The sales department of the organization, Book Sales, has these subdepartments:

- Wholesale
- Retail
- Online

To monitor the storage resources that the Books Sales department uses and that each of its subdepartments use, complete these tasks:

- Create a department for Book Sales.
- Create the Wholesale, Retail, and Online subdepartments that you want to add to the department:
 - Create the subdepartments as departments.
 - Add the departments as subdepartments to the department.
 - Add the applications to the subdepartments.

You can now view charts and information about the capacity and space usage of the department and each of the subdepartments on the Departments page and the Subdepartments page. You can also view charts and information about the performance of the storage resources that are associated with the applications or application subcomponents that the departments and subdepartments use.

Tip: In the **Subdepartments** column for the department on the Departments page, the name of the subdepartment is shown. If you add two or more departments to a subdepartment, the number of subdepartments is shown. To view information about the subdepartments, click the name of the subdepartment or the number.

You can view the total capacity and space usage for the department and the capacity and space usage for each of the subdepartments. You can add to your department hierarchy by adding more applications and departments.

Creating departments

To model the storage capacity that is consumed in a department for your business environment, in accordance with other department and application members, create departments for monitoring in IBM Spectrum Control.

Use the web-based GUI to create departments, subdepartments (member departments) and add applications for monitoring different levels of a business hierarchy.

To create a department, complete these steps:

1. In the menu bar, go to **Groups > Departments**.

- 2. On the Departments page, click Create Department.
- 3. On the **Create Department** page, enter information about the department.

You cannot have a duplicate department name on the same level in your department hierarchy. For example, you create a department named *Finance* and create a subdepartment named *Business*. You can create another department named *Business* as a top level department, but if you want to make your new *Business* department a subdepartment of the *Finance* department, you see a message that indicates there is already a department named *Business* that is a member of the *Finance* department; your department creation is denied.

You can customize the icon for your department by clicking the existing icon on the **Create Department** page and selecting another icon from the **Customize Icon** page.

- 4. Click Create and then Close.
- 5. Optional: On the Add Members page, click Add one or more applications.
 - a) On the populate page, select the applications that to add to your department.
 If you do not want to add any of the applications that are listed, click Create Application to create a new application to add to the department.
 - b) Click Finish.
- 6. Optional: On the Add Members page, click Create subdepartment .
 - a) On the Create Subdepartment page, enter your information.
 - b) Click Create.

To view the newly created department, go to the **Departments** page.

To see the applications that you added to your department, see the details page and click **Applications**.

To see the subdepartment you created, see the Subdepartments column on the **Departments** page or see the details page and click **Subdepartments**.

Continue to add to your department hierarchy by adding applications and departments.

Related tasks

"Adding and creating subdepartments" on page 408

To further enhance your data modeling at the department level and to add to your overall business hierarchy for monitoring, add an existing subdepartment or create a new subdepartment.

"Adding applications to departments" on page 409

To further enhance the data storage modeling in a business environment and see the storage capacity that is used at the different levels of your business hierarchy, add applications to departments. The departments you select from the **Add to Department** page become parents of the selected applications.

Adding and creating subdepartments

To further enhance your data modeling at the department level and to add to your overall business hierarchy for monitoring, add an existing subdepartment or create a new subdepartment.

To add a subdepartment to the department level hierarchy, use these steps:

- 1. In the menu bar, go to **Groups > Departments**.
- 2. On the **Departments** page, select one or more departments, right-click and click **Add to Department**.
- 3. On the **Add to Department** page, select one or more departments from the list and click **Save**.

 The department or departments you select from the **Add to Department** page become parents of the

selected department. The Hierarchy column shows the longest ancestry path of the department. You can also view the applications that are members of the department on the **Departments** page.

- 4. Optional: Create a new subdepartment:
 - a) Click Create Department to open the Create Department page.
 - b) Enter information about the department.
 - c) Click Create.
 - d) On the **Add Members** page, click **Create a subdepartment**.

- e) On the Create Subdepartment page, enter your information.
- f) Click Create.
- 5. Optional: Create a new subdepartment from the details page of the existing department:
 - a) In the **General** section, click **Subdepartments** to open the **Subdepartments** page.
 - a) Click Create Subdepartment to open the Create Subdepartment page.
 - b) Enter information about the subdepartment.
 - c) Click Create.

To see the subdepartment you created, see the Subdepartments column on the **Departments** page or click **Subdepartments** on the details page.

You can continue to create or add subdepartments to the department or add applications to the department.

Related tasks

"Creating departments" on page 407

To model the storage capacity that is consumed in a department for your business environment, in accordance with other department and application members, create departments for monitoring in IBM Spectrum Control.

"Adding applications to departments" on page 409

To further enhance the data storage modeling in a business environment and see the storage capacity that is used at the different levels of your business hierarchy, add applications to departments. The departments you select from the **Add to Department** page become parents of the selected applications.

Adding applications to departments

To further enhance the data storage modeling in a business environment and see the storage capacity that is used at the different levels of your business hierarchy, add applications to departments. The departments you select from the **Add to Department** page become parents of the selected applications.

To add applications to departments, complete these steps:

- 1. In the menu bar, go to **Groups** > **Applications**.
- 2. On the **Applications** page, select one or more applications, right-click and then, click **Add to..** > **Department**.
- 3. On the **Add to Department** page, select one or more departments to add to the selected application or applications.

The Hierarchy column shows the longest ancestry path of the department.

- 4. Click **Save**.
- 5. Optional: On the **Create Department** page, enter information about the department.
- 6. Click Create.
- 7. On the Add Members page, click Add one or more applications.
- 8. On the populate page, select the applications that to add to your department.

If you do not want to add any of the applications that are listed, click **Create Application** to create a new application to add to the department.

9. Click Finish.

View the added department or departments in the Departments column on the Applications page or the Departments tab of the application properties.

Related tasks

"Creating applications" on page 400

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies for applications, you can create applications for monitoring in IBM Spectrum Control.

"Creating departments" on page 407

To model the storage capacity that is consumed in a department for your business environment, in accordance with other department and application members, create departments for monitoring in IBM Spectrum Control.

Viewing information about departments

You can view detailed information about departments for use in storage capacity monitoring and performance troubleshooting. You can view additional information on the applications and subdepartments that are used to establish a hierarchal business structure.

To view the details of a department, use these steps:

- 1. In the menu bar, go to **Groups** > **Departments**.
- 2. On the **Departments** page, right-click a department in the list and select **View Details**.

The department details page shows information about the specific department. The information is available in the following sections on the page:

- **Overview**: The overall capacity trending and performance information about the department. An interactive breadcrumb that shows the hierarchy of the departments and applications and a list of all the top level departments.
- **Properties**: The key information about the department.
- **Applications**: The applications that belong to the department. If there is more than one application, the number of applications is displayed.
- **Subdepartments**: The subdepartments that belong to the department. If there is more than one subdepartment, the number of subdepartments is displayed.
- **Related Resources**: Information about related resources in the storage environment that the applications, which belong to the department, are using or are connected to.

Tip: To view descriptions of the information that is available on a department details page, click the **Help** icon **1** in the upper-right corner of the page.

Related concepts

Application and department hierarchies

The interactive breadcrumb on the Overview page provides a visual representation of the hierarchical relationships for applications and departments and enables you to navigate quickly between the elements in the hierarchy.

Related tasks

Creating a department hierarchy

To monitor the capacity and space usage of the department, you want to create a department hierarchy and associate applications with the department hierarchy. You can also monitor the performance of the storage resources that are associated with the applications that you add to the department hierarchy.

Removing departments and subdepartments

You can remove departments, subdepartments and applications that are members of the department. The data that is collected about the departments, subdepartments and applications is removed from the database repository.

To remove departments, subdepartments, and applications, that are members of the department, from a level of the business hierarchy, complete these steps:

- 1. In the menu bar, go to **Groups > Departments**.
- 2. On the **Departments** page, select one or more departments, right-click and click **Delete**.
- 3. Confirm the information and click **Delete**.
 - a) If you select **Delete subdepartments and applications**, the subdepartments and applications that belong only to the department, that is being removed, are removed.
 - b) If you do not select **Delete subdepartments and applications**, the subdepartments and applications of the department that are being removed become members of the parent department. If the department is a top level department, and the subdepartments are not members of any other department, they become top level departments.

The departments, subdepartments and applications are removed from the business level hierarchy and are no longer displayed on the **Departments** page.

To remove subdepartments from within the direct context of a department, use these steps:

- 1. In the menu bar, go to **Groups** > **Departments**.
- 2. On the **Departments** page, select a department from the list and right-click **View Details**.
- 3. Optional: Select a department from the **Departments** page and right-click **View Properties**.
 - a. Click the **Subdepartments** tab and continue with steps 5 and 7.
- 4. On the department details page, click **Subdepartments** in the General section.
- 5. On the **Subdepartments** page, select the subdepartment and right-click **Remove from Department**.
- 6. Optional: On the **Subdepartments** page, select the subdepartment and click **Delete**.
- 7. Reply to the confirmation message and click **OK**.

Removing applications from a department

You can remove an application from within a department. This action is only available from the department detail page or the properties for the department.

To remove applications from within the context of a department, use these steps:

- 1. In the menu bar, go to **Groups** > **Departments**.
- 2. On the **Departments** page, select a department from the list and right-click **View Details**.
- 3. Optional: Select a department from the **Departments** page and right-click **View Properties**.
 - a) Click the **Applications** tab and continue with steps 5-7.
- 4. On the department details page, click **Applications** from the General section.
- 5. On the **Applications** page, select the application and right-click **Remove from Department**.
- 6. Confirm the information.
- 7. Click OK.

The applications that were added within the context of a department are eligible for removal from within department.

Related tasks

"Viewing information about departments" on page 410

You can view detailed information about departments for use in storage capacity monitoring and performance troubleshooting. You can view additional information on the applications and subdepartments that are used to establish a hierarchal business structure.

Application and department hierarchies

The interactive breadcrumb on the Overview page provides a visual representation of the hierarchical relationships for applications and departments and enables you to navigate quickly between the elements in the hierarchy.

The breadcrumb uses default or customized icons to represent the related applications and departments. The greater-than sign (>) serves as the hierarchy separator and ends with the name of the selected application or department. Only the ancestor names of the selected application are shown.



When you move your mouse pointer over an icon in the breadcrumb, you see the siblings for that application or department. If you move your mouse pointer, for example, over the Data subcomponent, the related application and department elements are shown.

The visual representation of the hierarchy stops when either there are no additional parents or multiple parents are encountered from the selected application or department. When multiple parents are encountered, the number of parents are shown in parenthesis. For example, *Multiple* (2). Moving the mouse pointer over *Multiple* shows a list of the parents. The siblings of multiple parents are *not* shown.

Move your mouse pointer over a sibling in the breadcrumb and it becomes a name link. Click the application or department name link to navigate to the details for that application or department and view the breadcrumb if it exists.

You can add resources directly on the Overview page for applications from the breadcrumb. For example, in this case, you create an application named DB2 and create two subcomponents, Logs and Data. On the Overview page for the Logs subcomponent, click the Logs subcomponent name link in the breadcrumb and click **Filters** to create a resource filter for that subcomponent. Then, click the Data subcomponent name link and create a resource filter for that application subcomponent. To see a visual representation of the hierarchy for applications, subcomponents, departments, and subdepartments click **View Details**.

Related tasks

"Creating an application hierarchy" on page 399

To monitor, the capacity, space usage, and performance of the storage resources that your applications use, create application models.

"Creating a department hierarchy" on page 407

To monitor the capacity and space usage of the department, you want to create a department hierarchy and associate applications with the department hierarchy. You can also monitor the performance of the storage resources that are associated with the applications that you add to the department hierarchy.

Exporting information about applications and departments

You can save information about applications and departments to a PDF, CSV, or HTML file. Information that you can export from the GUI includes all the values that are being shown in the columns for a list of applications and departments.

The information that you export to a file is organized according to the sorting, filtering, and column order that is defined for a list. Before you export information, complete the following tasks to configure a list:

- Change the order of columns in the list as you want them to appear in the generated file.
- Hide columns in the list that you do not want to include in the generated file. Information in hidden columns is not exported.
- Sort the rows in the list as you want them to appear in the generated file. Each list in the GUI has a column or set of columns that determines the order of its rows.
- Filter the list to show only the applications and departments that you want to export. Use filtering to limit the rows that are shown in a list based on value in a specific column or set of default columns.

You can export information that is shown on the application and department list and detail pages.

Applications and departments list and detail pages

You can export information about the top level applications and departments and the related resources that are shown on application and department list and detail pages.

For example, you can export information about monitored applications and their subcomponents and departments on the Applications page. When you select the export action, the generated file includes the column values for each monitored application in the list.

For example, you can export information about the servers that are associated with an application on the details page.

- 1. In the GUI, go to the Applications or Departments page and select applications or departments for which you want to export information.
- 2. Click **Actions** and select **Export** > *file_format*, where *file_format* represents the format that you want to use for a file.

The following formats are available:

CSV (comma-separated values)

A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application. The information in a CSV file has the following format:

```
"Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","Column_name","C
```

where Column_name represents the name of a column in a table and data represents the data that is associated with a column.

PDF

You can view a PDF file with Acrobat reader. The information that you export to a PDF file is formatted into a table.

Tip: When you export a table of data that contains many columns, the rows in that table might span multiple pages in the resulting PDF. For example, if a table contains 20 columns, then the row for a specific resource might be shown on more than one page (10 columns on the first page and 10 columns on the second page).

To reduce the number of pages that a table spans, before you select the export action, hide the columns that you do not want to include in the PDF. Continue hiding columns until the table no longer spans multiple pages.

HTML

You can view an HTML file with a web browser. The information that you export to an HTML file is formatted into a table.

- 3. Optional: Depending on how your web browser is configured, you can specify the name and location of the generated file.
- 4. Click the save option to export the information to a file.

Related concepts

"Customizing lists" on page 454

Customize lists to focus on the information that is important to you. You can filter information, sort rows, and show, hide, and reorder columns in the table views for resources, tasks, alerts, and other objects.

General groups

View and administer the general groups that are used to alert on logically related storage resources. You can define alerts for a general group to notify you about changes in the configuration, attributes, and performance of the resources in the group.

Organizing your resources into general groups can be helpful in the following situations:

• When you want to receive alert notifications about changes for a subset of the resources of a particular type. For example, you want to detect when the ports that are used for replication on your SAN Volume Controller have insufficient buffer-to-buffer credit. You do not want the alerts to apply to other ports on

your SAN Volume Controller. You can group the ports that are used for replication and then define alerts for the group.

• When you want to receive alert notifications about changes for a group of resources that are logically related. Examples include all the servers that use a particular operating system or all the storage systems at a specific location.

General group hierarchies

Organizing resources into general groups and their subgroups can be helpful when you want to quickly view information about a group of resources, but you also want to view information about subgroups of resources within the group.

Only top-level groups in the general group hierarchy are shown on the **General Groups** page. If the general group has a subgroup, the name of the subgroup is shown in the Subgroups column. If the general group has more than one subgroup, click the number of groups in the Subgroups column to view the subgroups.

Related tasks

Defining alerts for general groups

Define alerts for changes in the configuration, attributes, and performance of the resources in your general groups.

Creating a general group hierarchy

Organizing resources into general groups and their subgroups can be helpful when you want to quickly view information about a group of resources, but you also want to view information about subgroups of resources within the group.

In this scenario, you want to monitor all the ports on your SAN Volume Controller but you also want to separately monitor the following subsets of ports on the SAN Volume Controller:

- Ports that are used for inter-node communication
- Ports that are used for host I/O exchanges

To add the ports to a general group hierarchy, complete the following steps:

- 1. Add the ports that are used for inter-node communication to a general group, SVC Inter Node Ports.
- 2. Add the ports that are used for host I/O exchanges to another general group, SVC Host I/O Ports.
- 3. Add the groups as subgroups of a parent general group, SVC All Ports.

You can now quickly view information about the resources in the hierarchy on the details page for the general groups.

When you view the SVC Inter Node Ports or SVC Host I/O Ports groups, you see information about the specific ports that are members of that subgroup.

When you view the parent group, SVC All Ports, you see information about the ports that are members of the parent group and also the ports in the subgroups.

Creating general groups and adding resources

When you create and associate storage resources with a general group, you can quickly view information about the group and its resources.

You can add the following resources and their internal resources to a general group:

- Storage systems
- Servers
- Hypervisors
- Switches
- Fabrics
- · Other general groups

Only resources that you specifically add to a group are included as members of the group. For example, if you add a SAN Volume Controller, the internal resources of the SAN Volume Controller, such as volumes and pools, are not automatically added to the group. To add the volumes and pools, you must specifically select those resources and add them to the group.

To create a general group and add resources to the group, complete these steps:

- 1. Take one of the following actions to go to the list page for the resource or group that you want to add:
 - To add top-level resources such as storage systems, servers, hypervisors, switches, or fabrics, go
 to the appropriate resource page. For example, to add servers, in the menu bar, click Servers >
 Servers.
 - To add internal resources of top-level resources, complete the following steps:
 - a. Go to the resource list page for the top-level resource. For example, to add internal resources of a SAN Volume Controller, in the menu bar, click **Storage > Block Storage Systems**.
 - b. Right-click the resource and click **View Details**.
 - c. In the Internal Resources section, click the type of resource that you want to add, for example, volumes or pools.
 - To add general groups, in the menu bar, click **Groups** > **General Groups**.
- 2. Right-click one or more resources or groups and click **Add to General Group**.
- 3. Click **Add to new group** and specify a name and description for the new group.
- 4. Optional: Customize the icon for a new group by clicking the existing icon and selecting another icon.
- 5. Click Save.

You can now quickly view information about the resources in the general group hierarchy on the details page for the general group.

Related tasks

Adding resources to general groups

Add resources to general groups so that you can receive alert notifications about changes in the configuration, attributes, and performance of the resources in the group. You can add resources to one or more existing groups or you can create a new group.

Related reference

Viewing and administering general groups

View and administer the general groups that you use to alert on logically related storage resources.

Adding general groups as subgroups

When you add one or more general groups as subgroups to another general group, you can quickly view information about the resources in the group hierarchy.

- 1. In the menu bar, click **Groups** > **General Groups**.
- 2. Right-click one or more general groups and click **Add to General Group**.
- 3. Take one of the following actions:
 - To add the groups as subgroups of existing groups, click **Add to existing groups**, click one or more groups in the list, and click **Save**.
 - To add the groups as subgroups of a new group, click **Add to new group**, specify a name and description for the new group, and click **Save**.

Related tasks

Adding resources to general groups

Add resources to general groups so that you can receive alert notifications about changes in the configuration, attributes, and performance of the resources in the group. You can add resources to one or more existing groups or you can create a new group.

Related reference

Viewing and administering general groups

View and administer the general groups that you use to alert on logically related storage resources.

Adding resources to general groups

Add resources to general groups so that you can receive alert notifications about changes in the configuration, attributes, and performance of the resources in the group. You can add resources to one or more existing groups or you can create a new group.

The following resources and their internal resources can be added to a general group:

- Storage systems*
- Servers
- Hypervisors
- Switches
- Fabrics
- · Other general groups

Tip: * For IBM Cloud Object Storage, only vaults can be added to general groups.

Only resources that you specifically add to a group are included as members of the group. For example, if you add a SAN Volume Controller, the internal resources of the SAN Volume Controller, such as volumes and pools, are not automatically added to the group. To add the volumes and pools, you must specifically select those resources and add them to the group.

You can create a general group hierarchy by adding one or more general groups as subgroups of a general group. When you define a general group alert for a resource such as a volume, the alert applies to all the volumes that belong to that group and all its subgroups.

To add resources to a general group, complete these steps:

- 1. Take one of the following actions to go to the list page for the resource or group that you want to add:
 - To add top-level resources such as storage systems, servers, hypervisors, switches, or fabrics, go
 to the appropriate resource page. For example, to add servers, in the menu bar, click Servers >
 Servers.
 - To add internal resources of top-level resources, complete the following steps:
 - a. Go to the resource list page for the top-level resource. For example, to add internal resources of a SAN Volume Controller, in the menu bar, click **Storage > Block Storage Systems**.
 - b. Right-click the resource and click View Details.
 - c. In the Internal Resources section, click the type of resource that you want to add, for example, volumes or pools.
 - To add general groups, in the menu bar, click **Groups** > **General Groups**.
- 2. Right-click one or more resources or groups and click **Add to General Group**.
- 3. Take one of the following actions:
 - To add the resources or groups to a new group, click **Add to new group** and specify a name and description for the new group.
 - To add the resources or groups to one or more existing groups, click **Add to existing groups** and click the appropriate groups in the list.
- 4. Optional: Customize the icon for a new group by clicking the existing icon and selecting another icon.
- 5. Click **Save**.

To view the resources and subgroups that are members of the group, complete the following steps:

- 1. Go to Groups > General Groups.
- 2. Right-click the general group and select View Details.

On the general group details page, all the different resource types, such as volumes or pools, that are members of the group or members of its subgroups are shown in the Members section.

- 3. Click the resource type to view the specific resources. For example, click **Volumes** to view all the volumes that are in the general group or in its subgroups.
- 4. Click **Subgroups** to view the child groups.

To define alerts for the general group, use the View Alert Definitions action on the General Groups page.

Related tasks

Creating general groups and adding resources

When you create and associate storage resources with a general group, you can quickly view information about the group and its resources.

Adding general groups as subgroups

When you add one or more general groups as subgroups to another general group, you can quickly view information about the resources in the group hierarchy.

Related reference

Viewing and administering general groups

View and administer the general groups that you use to alert on logically related storage resources.

Viewing and administering general groups

View and administer the general groups that you use to alert on logically related storage resources.

Administering general groups

Table 217. Administering general groups	
Actions	Navigation
Create a general group and add resources.	 In the menu bar, click the resource list page for the resources that you want to add. For example, if you want to add file storage systems, click Storage > File Storage Systems. Right-click one or more resources and click Add to General Group. Click Add to new group, specify a name and description for the
	new group, and click Save .
Add resources to general groups.	In the menu bar, click the resource list page for the resources that you want to add. For example, if you want to add block storage systems, click Storage > Block Storage Systems.
	Right-click one or more resources and click Add to General Group .
	3. Take one of the following actions:
	 a. To add the resources to one or more existing groups, click Add to existing groups, click the appropriate groups in the list, and click Save.
	 b. To add the resources to a new group, click Add to new group, specify a name and description for the new group, and click Save.

Table 217. Administering general groups (continued)	
Actions	Navigation
Set which alert policy manages a general group <u>Learn more</u>	 Go to the details page for the general group for which you want to set the policy. In the menu bar, click Groups > General Groups. Right-click one of the general groups, then click View Details. Click Alerts in the General section. Click Set Policy from the Policy Actions menu. To set the alert policy for multiple general groups, go to Settings > Alert Policies. Double-click the policy, click the resources tab, then click Edit Resources.
Create an alert policy from the alert definitions and notification settings in a general group	 Go to the details page for the general group from which you want to create the policy. In the menu bar, click Groups > General Groups. Right-click one of the general groups, then click View Details. Click Alerts Definitions in the General section. Click Create Policy from the Policy Actions menu.
View or modify the alert policy that manages a general group Learn more about defining alerts and notification settings	 Go to the details page for the general group whose policy you want to view. In the menu bar, click Groups > General Groups. Right-click one of the general groups, then click View Details. Click Alerts in the General section. Click View Policy from the Policy Actions menu.
Delete general groups.	 In the menu bar, click Groups > General Groups. Right-click one or more general groups and click Delete. To delete subgroups, ensure that the Delete subgroups? check box is selected. When a group is deleted, the information that is associated with the group, such as alerts and alert definitions, is removed from the product. The resources that were members of the group are removed from the group but are still monitored by the product.
Remove resources from a general group.	 In the menu bar, click Groups > General Groups. Right-click a general group and click View Details. In the Members section, click the type of resource that you want to remove. For example, if you want to remove some of the volumes from the group, click Volumes. Right-click the resources that you want to remove and click Remove from General Group. Click OK to confirm that you want to remove the resources from the general group and from all of its subgroups.

Table 217. Administering general groups (continued)	
Actions	Navigation
Add general groups as subgroups.	1. In the menu bar, click Groups > General Groups .
	2. Right-click one or more general groups and click Add to General Group .
	3. Take one of the following actions:
	a. To add the groups as subgroups of existing groups, click Add to existing groups , click one or more groups in the list, and click Save .
	b. To add the groups as subgroups of a new group, click Add to new group , specify a name and description for the new group, and click Save .
Remove subgroups from a group	1. In the menu bar, click Groups > General Groups .
hierarchy.	2. Right-click a general group and click View Details .
	3. Click Subgroups in the General section.
	4. Right-click the general groups that you want to remove and click Remove from General Group .
	When you remove a subgroup from its parent group, the subgroup is moved to the same level in the hierarchy as the parent group. The subgroup is still monitored by the product.
Create a subgroup.	1. In the menu bar, click Groups > General Groups .
	Right-click the general group that you want to create a subgroup for and click View Details .
	3. Click Subgroups in the General section.
	4. Click Create Subgroup .
	5. Specify a name and description for the new group.
	6. Customize the icon for the group by clicking the existing icon and selecting another icon.
	7. Click Create .
Delete subgroups.	1. In the menu bar, click Groups > General Groups .
J .	Right-click the general group that contains the subgroups that you want to delete and click View Details .
	3. Click Subgroups in the General section.
	4. Right-click one or more subgroups and click Delete .
	5. To delete further subgroups, ensure that the Delete subgroups? check box is selected.
	When a group is deleted, the information that is associated with the group, such as alerts and alert definitions, is removed from the product. The resources that were members of the group are removed from the group but are still monitored by the product.

Viewing general groups

Table 218. Viewing general groups	
Actions	Navigation
View a list of general groups.	In the menu bar, click Groups > General Groups .
	The following general groups are shown:
	The top-level groups in general group hierarchies
	All general groups that are not part of a general group hierarchy
View subgroups in a general	1. In the menu bar, click Groups > General Groups .
group hierarchy.	Click the number of groups or the group name that is shown in the subgroups column.
	3. Click the Subgroups tab.
View details about a general group, including the alert notifications, the resources that are members of the group, and the subgroups.	In the menu bar, click Groups > General Groups . Right-click a general group and select View Details .
View the alert definitions and	1. In the menu bar, click Groups > General Groups .
notification settings for a general group.	Right-click a general group and select View Alert Definitions.
View the alerts that were	1. In the menu bar, click Groups > General Groups .
triggered for a general group.	2. Right-click a general group and select View Details .
	3. Click Alerts in the General section.
View reports about general groups.	In the menu bar, click Cognos . The Cognos Analytics reporting tool is included in Tivoli Common Reporting.
	To view reports about general groups, view reports about storage resource groups in the Cognos Analytics reporting tool.

Related tasks

Creating general groups and adding resources

When you create and associate storage resources with a general group, you can quickly view information about the group and its resources.

Adding resources to general groups

Add resources to general groups so that you can receive alert notifications about changes in the configuration, attributes, and performance of the resources in the group. You can add resources to one or more existing groups or you can create a new group.

Adding general groups as subgroups

When you add one or more general groups as subgroups to another general group, you can quickly view information about the resources in the group hierarchy.

Provisioning storage

You can assign storage to servers, hypervisors, and clusters on servers and hypervisors in the GUI.

Discontinued support: Provisioning is no longer supported in IBM Spectrum Control. While the feature might still work in this release, it's recommended that you use another tool for your provisioning needs,

when possible. For a complete list of discontinued features, see "Discontinued features in IBM Spectrum Control" on page xxx.

You can also assign the following types of storage to servers, hypervisors, and clusters on servers and hypervisors:

- Volumes
- · NAS file shares

When you assign storage, you must associate the storage with a service class. You can associate the storage with a default service class such as gold, silver, or bronze, or you can create a user-defined service class. You can also create capacity pools to restrict the provisioning of storage to a defined set of storage resources such as pools on a storage system in a particular location.

Configuring IBM Spectrum Control for provisioning

Before you can use the simplified provisioning capabilities that are provided with IBM Spectrum Control, you must configure your environment.

Discontinued support: Cloud configuration, provisioning, and optimization are no longer supported in IBM Spectrum Control. While these features might still work in this release, it's recommended that you use another tool for your provisioning and optimization needs, when possible. For a complete list of discontinued features, see "Discontinued features in IBM Spectrum Control" on page xxx.

To provision storage, you are required to specify only the storage capacity and storage quality that is required. After volumes are created, IBM Spectrum Control can analyze and optimize volume performance.

To configure IBM Spectrum Control for provisioning, complete the following tasks:

- 1. Assign storage pools to tiers by setting the tier level of each storage pool. The tier levels are number tags that can reflect any tier structure in the environment.
 - You must assign storage pools to tiers to later provision volumes that require a certain tier level and to optimize the placement of volumes. Storage pools do not need to be assigned to tiers to provision volumes that do not require a tier level, or to provision file shares.
- 2. Create service classes. A service class is a set of properties that describe capabilities and characteristics of storage resources. A service class typically describes a particular quality of service, and is used during provisioning to describe storage requirements.
 For example, a block-storage service class specifies properties such as a required RAID level, and whether storage resources must be able to encrypt or thin provision volumes.
- 3. Optional: Create capacity pools. You can restrict provisioning requests to a capacity pool. Capacity pools are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business.

 For example, a capacity pool might contain the storage resources that are allocated to a particular department or division of your business, or to a particular application.

Configuring service classes

Configure service classes to represent the different levels or types of storage quality that you want to provision. When you are configuring service classes for your installation, consider whether certain storage requests must always be satisfied from a particular set of resources. Also, consider whether you want to allow users without administrator privileges to provision storage, and whether you want their provisioning requests to require administrator approval.

Discontinued support: Cloud configuration, provisioning, and optimization are no longer supported in IBM Spectrum Control. While these features might still work in this release, it's recommended that you use another tool for your provisioning and optimization needs, when possible. For a complete list of discontinued features, see "Discontinued features in IBM Spectrum Control" on page xxx.

Creating service classes

Create service classes to specify storage requirements for future provisioning requests. A service class describes attributes and capabilities of block storage or file storage resources, and typically describes particular quality of service.

You must have Administrator privileges to create a service class.

The IBM Spectrum Control GUI guides you through the steps to create a service class. When you provision a volume or share, you define storage requirements by specifying a service class. A block-storage service class specifies the requirements for the block-storage resource from which a new volume is allocated. A file-storage service class defines requirements for the file storage resources from which a new share is allocated.

You can associate a service class with one or more capacity pools. If you do, future provisioning requests can be satisfied only by storage resources in the specified capacity pools.

You can selectively grant users who do not have administrator privileges permission to provision storage by using the service class.

Restrictions:

- To create a block-storage service class, you must have the IBM Spectrum Control Advanced Edition license.
- To associate the service class with a capacity pool, one or more capacity pools must be defined in the IBM Spectrum Control database. If no capacity pool is defined, the Specify Capacity Pool page is not displayed.

Start this task at the **Service Classes** page. To open the Service Classes page, select **Advanced Analytics** > **Cloud Configuration** > **Work With Service Classes**. To create a service class, complete the following steps:

- 1. From the Service Classes page, click Create Service Class.
- 2. Select the type of service class that you want to create:
 - To create a block-storage service class, click the **Block** icon.
 - To create a file-storage service class, click the **File** icon.
- 3. In the **Define Properties** page, complete the following steps:
 - a) Specify a name for the service class.
 - b) Optional: Specify a service class description.
 - Descriptions are displayed in the table of service classes on the Service Classes page.
 - c) Specify the general attributes of the service class.
 - Most of the general properties specify requirements for provisioning. Some of the general properties specify configuration instructions for the storage resources or for the volumes or shares that are provisioned from them.

Tip: To display help information about a service class property, move the mouse pointer over the field or control. Then, move the mouse pointer over the question mark icon displayed next to the field.

- 4. Optional: Specify or modify the advanced properties of the service class.
 - a) From the **Define Properties** page, click **Advanced**.
 - b) If you are creating a block-storage service class, and thin provisioning is enabled for the service class, specify configuration instructions for thin provisioning.
 - c) If you are creating a block-storage service class, select how the multipath driver on the host is configured. For hypervisors, the multipath driver is configured only if provisioning is initiated from the vSphere Web Client extension for IBM Spectrum Control.
 - d) Create custom requirements for the service class by specifying up to three custom tags. To provide the service class, storage resources must have all the same tags that are specified in the service class.

- 5. Optional: In the **Specify Capacity Pool** page, associate the service class with one or more capacity pools.
 - If capacity pools are associated with the service class, future provisioning requests that specify the service class can be satisfied only by storage resources in the capacity pools.
- 6. Optional: Grant permission to individual users, who do not have administrator privileges, to provision storage by using the service class.
 - By default, users must be assigned to the Administrator role to provision storage. However, you can selectively grant users who are limited to the Monitor or External Application role permission to provision by using this service class. You can also specify whether administrator approval is required to run provisioning tasks that are created by these users.

7. Click Finish

The new service class is listed in the table of service classes on the **Service Classes** page, and can be specified in future provisioning requests.

Viewing the details of service classes

You can view the provisioning requirements and storage constraints that a service class defines. You can also view the volumes or shares that were provisioned by using the service class.

From the **Service Classes** page, you can open a properties page to view the following information for a service class:

- The properties of a service class, such as the required RAID level or storage tier for a block-storage service class.
- The storage constraints that are imposed by the service class. The service class can allow provisioning from all available storage, or can limit provisioning to one or more capacity pools.
- The users who can provision storage by using the service class.
- The volumes or shares that were provisioned by using the service class.
- The candidate storage for the service class. The *candidate storage* is the set of storage resources that can provide the service class.

To open the service class properties window from the **Service Classes** page, complete the following steps:

1. In the menu bar in the web-based GUI, go to **Advanced Analytics** > **Cloud Configuration**, and click **Work With Service Classes**.

The **Service Classes** page is displayed.

From the list of service classes, select the service class, and then select Actions > View / Modify.
 The service class properties window is displayed.

Modifying service classes

You can modify the attributes of a service class to change the provisioning requirements, capacity pool storage constraints, user permissions, and service class name.

You must have Administrator privileges to modify a service class.

Important: When volumes or shares are provisioned by specifying a service class, they are associated with the service class. The **Volumes** page and the **Shares** page in the web-based GUI can display the name of the service class that was used to provision the volume or share. If you modify a service class, the volumes or shares are still associated with the service class, but might be on storage resources that no longer satisfy the requirements of the service class. Depending on your changes to the service class, users might incorrectly assume that the new volumes or shares have attributes that they do not possess.

Start this task at the **Service Classes** page. To open the Service Classes page, select **Advanced Analytics** > **Cloud Configuration** > **Work With Service Classes**. To modify a service class, complete the following steps:

- From the list of service classes, select a service class and then select Actions > View / Modify.
 The service class properties window is displayed.
- 2. Optional: Click the **General** tab to modify the general attributes of the service class.

Most of the general properties specify requirements for provisioning. Some of the general properties specify configuration instructions for the storage resources or for the volumes or shares that are provisioned from them.

Tip: To display help information about a service class property, move the mouse pointer over the field or control. Then, move the mouse pointer over the question mark icon displayed next to the field.

- 3. Optional: Specify or modify the advanced properties of the service class.
 - a) From the General tab, click Advanced.
 - b) If you are modifying a block-storage service class, and thin provisioning is enabled for the service class, specify configuration instructions for thin provisioning.
 - c) If you are modifying a block-storage service class, select how the multipath driver on the host is configured.
 - d) Create custom requirements for the service class by specifying up to three custom tags. To provide the service class, storage resources must have all the same tags that are specified in the service class.
- 4. Optional: Modify storage constraints.

If one or more capacity pools are defined in the IBM Spectrum Control database, you can associate the service class with one or more capacity pools. If capacity pools are associated with the service class, future provisioning requests that specify the service class can be satisfied only by storage resources in the capacity pools. If a capacity pool was previously associated with the service class, you can disassociate it from the service class.

5. Optional: Modify user permissions.

By default, users must be assigned to the Administrator role to provision storage. However, you can selectively grant users who are limited to the Monitor or External Application role permission to provision by using this service class. You can also specify whether administrator approval is required to run provisioning tasks that are created by these users.

If users in the Monitor or External Application roles have permission to provision by using the service class, you can revoke that permission. You can specify that only administrators can provision by using the service class, or you can selectively remove users. To remove a user, right-click the user name in the table and then select **Remove** from the context-sensitive menu.

Deleting service classes

You can delete a service class that is no longer needed.

You must have Administrator privileges to delete a service class.

Deleting a service class does not affect the storage volumes or NAS shares that are already provisioned by using the service class. However, the storage volumes or NAS shares are no longer associated with any service class. When a volume that is not associated with a service class is optimized, the volume can be migrated to any available storage pool. The storage pool to which the volume is migrated might not satisfy the requirements of the service class that was specified when the volume was provisioned.

To delete a service class from the **Service Classes** page, complete these steps:

1. In the menu bar, go to Advanced Analytics > Cloud Configuration, and click Work With Service Classes.

The **Service Classes** page is displayed.

2. From the list of service classes, select the service class, and then select **Actions > Delete**.

Tagging resources to satisfy custom requirements

A service class can specify tags to create custom requirements for provisioning. Add the same tags to storage resources that satisfy the custom requirements. When the service class is specified during provisioning, only the tagged resources are candidates for provisioning.

You must have Administrator privileges to tag storage resources.

When you provision storage, you indicate your storage requirements by specifying a service class. Certain properties of the service class describe capabilities that storage resources must have to be a candidate

for provisioning. In addition to the default properties, a service class can define custom requirements by specifying up to three custom tags. To provide the service class, a storage resource must have all the same tags as the service class. The ability to define custom requirements by tagging service classes and storage resources gives you more control over the selection of storage resources during provisioning. By using tags, you can filter the candidates for provisioning by using any criteria that you choose.

When you provision volumes, you specify the requirements by using a block-storage service class. Based on the requirements of the service class, IBM Spectrum Control identifies a storage pool for the volume. If the service class specifies tags, only pools that have all the same tags are candidates for provisioning. If a pool is not tagged, any tags on the containing storage system also apply to the pool.

When you provision shares, you specify the requirements by using a file-storage service class. Based on the requirements of the service class, IBM Spectrum Control identifies a file system or Network Shared Disk (NSD) for the share. If the service class specifies tags, only file systems and NSDs that have all the same tags are candidates for provisioning. If a file system or NSD is not tagged, any tags on the containing storage system also apply to the internal resource.

Tag only the storage resources that satisfy the intended custom requirements of the service class. During provisioning, only the tagged resources can provide the service class.

1. Open the resource list page for the type of storage resource that you want to tag. To satisfy the custom tag requirements of a block-storage service class, you can tag a storage pool or its containing block storage system. To satisfy the custom tag requirements of a file-storage service class, you can tag a file system or NSD, or the containing file storage system.

Important: The tags on a containing storage system apply to an internal pool, file system, or NSD resource only if the internal resource is not tagged.

- If you are tagging a storage pool, from the menu bar, go to Storage > Pools.
- If you are tagging the block storage system that contains the pools, complete the following steps:
 - a. In the menu bar, go to **Storage > Block Storage Systems**.
 - b. Click the Block Storage tab.
- If you are tagging a file system or NSD, complete the following steps:
 - a. In the menu bar, go to Storage > Block Storage Systems.
 - b. Click the **File Storage** tab.
 - c. From the list of storage systems, select the storage system that contains the file system or NSD, and then select **View Details**. Alternatively, you can select **Actions** > **View Details**.
 - d. In the Internal Resources area, click File Systems or Network Shared Disks.
- If you are tagging the file storage system that contains the file systems or NSDs, complete the following steps:
 - a. In the menu bar, go to **Storage > Block Storage Systems**.
 - b. Click the **File Storage** tab.
- 2. In the table of storage systems, pools, file systems, or NSDs, select the storage resource that you want to tag, and then select **View Properties**. Alternatively, you can select **Actions** > **View Properties**.
- 3. In the properties notebook for the storage system, pool, file system, or NSD, click **Edit**.
- 4. In the custom tag fields of the properties notebook, specify tags that match the tags that are specified in the service class. To provide the service class, the resource must have all of the same tags.
- 5. In the properties notebook, click **Save**.

Configuring capacity pools

Configure capacity pools to track the used and available capacity for block and file storage on any set of storage resources. Provisioning requests can also be restricted to resources in a capacity pool.

Discontinued support: Cloud configuration, provisioning, and optimization are no longer supported in IBM Spectrum Control. While these features might still work in this release, it's recommended that you

use another tool for your provisioning and optimization needs, when possible. For a complete list of discontinued features, see "Discontinued features in IBM Spectrum Control" on page xxx.

Adding resources to capacity pools

Divide storage resources into capacity pools to later restrict provisioning or optimization to a specific set of resources. You can also use capacity pools to track the used and available capacity for block and file storage on any set of storage resources. You can add resources to an existing capacity pool, or you can create a new capacity pool.

You must have Administrator privileges to perform this task.

You can add storage resources to a capacity pool from resource list pages such as the Storage Systems page and the Pools page.

You can add the following resource types to a capacity pool:

- · Storage systems
- · Storage pools
- File systems of file storage systems

When you add a storage system to a capacity pool, any storage pool or file system that is an internal resource of that storage system is also indirectly assigned to the capacity pool.

Restriction: A resource can be assigned to only one capacity pool. When you attempt to add one or more resources to a capacity pool that are already assigned to different capacity pool, IBM Spectrum Control displays a message that lists the resources and their capacity pool assignments. When a message is displayed, ensure that you want to change the capacity pool assignments for the listed resources.

To add one or more resources to a capacity pool, complete these steps:

- 1. Open the appropriate list page for the resource you want to add. You can add an entire storage system, or you can add individual pools or file systems.
 - To add a storage system, in the menu bar, go to Storage > Block Storage Systems or Storage > File Storage Systems.
 - To add a storage pool, in the menu bar, go to **Storage** > **Pools**.
 - To add a file system, complete the following steps:
 - a. In the menu bar, go to Storage > File Storage Systems.
 - b. Select the storage system that contains the file system, then select **View Details** from the context-sensitive menu.
 - c. From the Internal Resources list, click File Systems or Network Shared Disks.
- 2. To add one or more resource to a capacity pool, select the resources in the list, and then select **Actions** > **Add to Capacity Pool**.

If one or more capacity pools are defined, a list of capacity pools is displayed. If no capacity pools are defined, the **Create Capacity Pool** window is displayed so that you can add resources to a new capacity pool.

Tip: The list of capacity pools and the **Create Capacity Pool** window shows capacity information. For the list of capacity pools, the Capacity column shows a bar that summarizes the used and free space in each capacity pool. In the **Create Capacity Pool** window, a capacity bar summarizes the used and free space of the resources that are being added to the new capacity pool. Move the mouse pointer over the bar to display a text summary and details of the capacity usage.

- 3. Take one of the following actions:
 - If the Create Capacity Pool window is displayed, go to step "4" on page 427.
 - If a list of capacity pools is displayed, add the selected resources to an existing capacity pool or to a new capacity pool.
 - To add the selected resource or resources to an existing capacity pool, select the appropriate row in the list of capacity pools, and click **Save**.

- To add the selected resource to a new capacity pool, click Create Capacity Pool. The Create Capacity Pool window is displayed.
- 4. In the **Create Capacity Pool** window, specify a name for the capacity pool.
- 5. Optional: Specify a capacity pool description.
 - Capacity pool descriptions appear in the list of capacity pools on the Capacity Pools page.
- 6. Optional: Specify custom tags to associate any significant information with the capacity pool.

 You can specify up to three custom tags. The custom tags can be displayed in the list of capacity pools on the Capacity Pools page.
- 7. Save the new capacity pool by clicking **OK**.

Viewing the details of capacity pools

You can view details of a capacity pool, such as its total capacity, used and available capacity, and member resources.

The **Capacity Pools** page lists all of the capacity pools that are defined in IBM Spectrum Control. For any capacity pool in the list, you can open a window to view more details about the capacity pool and its storage resources.

To view details of a capacity pool, complete these steps:

1. In the menu bar, go to Advanced Analytics > Cloud Configuration, and click Work With Capacity Pools.

The **Capacity Pools** page is displayed.

- 2. From the list of capacity pools, select the capacity pool, and then select **Actions** > **View / Modify**. A window is displayed, showing the following details:
 - A bar summarizing the used and free space in the capacity pool in gibibytes. The bar represents the total capacity of the capacity pool. The bar is divided into sections that represent the following measurements:
 - The amount of capacity that is used for file storage
 - The amount of available capacity for file storage
 - The amount of capacity that is used for block storage
 - The amount of block storage capacity that is not used for volumes, and is not reserved by pending or scheduled provisioning tasks.
 - The amount of block storage capacity that is reserved by pending or scheduled provisioning tasks

Move the mouse pointer over the bar to display a text summary and details of the capacity usage, and a legend for interpreting the colors on the bar.

A table that lists the storage systems, storage pools, and file systems that are in the capacity pool.
 To view details of any of these resources, select the resource, and then select Actions > View
 Properties.

Modifying capacity pools

You can modify the attributes of a capacity pool such as its name and description, and you can remove storage resources from a capacity pool.

You must have Administrator privileges to perform this task.

To modify a capacity pool, complete these steps:

1. In the menu bar in the web-based GUI, go to **Advanced Analytics** > **Cloud Configuration**, and click **Work With Capacity Pools**.

The Capacity Pools page is displayed.

- 2. From the list of capacity pools, select the capacity pool, and then select Actions > View / Modify. A window that shows details of the capacity pool is displayed. A table lists the storage systems, storage pools, and file systems that are in the capacity pool.
- 3. Optional: Modify attributes of the capacity pool.

4. Optional: Remove one or more resources from the capacity pool. Removing a storage resource from a capacity pool does not delete the storage resource, but does remove its capacity from the capacity pool. To remove resources, select the appropriate resources in the list, and then select **Actions** > **Remove**.

The resources are removed from the list, but are not removed from the capacity pool until you save your changes.

5. Click Save.

Deleting capacity pools

You can delete a capacity pool that is no longer needed.

You must have Administrator privileges to perform this task.

When you delete a capacity pool, the storage resources that were added to the capacity pool are no longer members of any capacity pool. Deleting a capacity pool does not affect any volumes or NAS shares that were provisioned from the capacity pool. However, the volumes or NAS shares are no longer associated with a capacity pool.

To delete a capacity pool, complete the following steps:

1. In the menu bar in the web-based GUI, go to **Advanced Analytics** > **Cloud Configuration**, and click **Work With Capacity Pools**.

The Capacity Pools page is displayed.

- 2. From the list of capacity pools, select the capacity pool, and then select Actions > Delete
- 3. In the confirmation window, click **Delete**.

Block storage: Calculating available capacity and determining the placement of volumes

For each block-storage service class, IBM Spectrum Control calculates the amount of available capacity for volumes of that service class. When you request volumes of a particular service class, IBM Spectrum Control recommends the best location for the volumes from the available capacity.

The **Service Classes** page in the GUI displays the available capacity for each block-storage service class. The available capacity is the amount of space that is available across storage pools for new volumes of the service class. Typically, not all pools are used in the available-space calculation. Some pools might be eliminated because of the consolidated status of the pool or its storage system. Other pools might be eliminated because they do not satisfy the requirements of the service class. The set of storage pools that are used to calculate the available capacity are also the candidates for provisioning by using the service class.

When you request storage volumes by using the Provision Storage wizard, you specify the capacity that is required for one or more volumes, and you specify a service class. You can also specify a capacity pool from which the volumes must be provisioned. Based on the volume size that is requested and if a capacity pool is specified, the candidates for provisioning are further refined to include only the pools that have enough available capacity and are in the specified capacity pool. From the remaining candidates for provisioning, IBM Spectrum Control creates a provisioning task that identifies the best location for the new volumes.

How available capacity for a block-storage service class is calculated

For each block-storage service class, IBM Spectrum Control calculates the available capacity for volumes of that service class. Only storage pools that are candidates for provisioning by using the service class are used to calculate the available capacity.

Identifying candidates for provisioning is a filtering process. For example, if the service class specifies that RAID 10 is required, then any storage pool that is not configured at RAID 10 is eliminated as a candidate for provisioning. IBM Spectrum Control filters the set of available storage resources against certain properties of the service class until only the storage resources that can satisfy all of the requirements remain. These remaining storage resources are the candidates for provisioning. Only certain properties of a service class are used to identify candidates for provisioning. More advanced properties contain configuration instructions, and are not used to identify candidates for provisioning.

The candidates for provisioning are identified by eliminating storage pools that cannot provide the service class, as follows:

- If a storage pool is a primordial pool or if it is formatted for count key data (CKD) volumes, it is not a candidate for provisioning.
- If a storage system has a condition of Error or Unreachable, none of its internal storage pools are candidates for provisioning. If a storage pool has a status of Error, it is not a candidate for provisioning.

Remember: Status is different from condition. Status is the status reported by a storage resource without considering the status of its internal resources. Condition is the combined status of the monitored, internal resources for the higher-level resource, and the status of the higher-level resource itself. The condition of a storage system is shown on the **Storage Systems** page and the condition for a pool is shown on the **Pools** page. The condition of a resource might not have the same value as the status, which is used to eliminate pools as candidates for provisioning.

Although the condition of a resource might not have the same value as its status, the condition is always at least as severe as the status. The condition reflects the most severe status of any internal resource or the status of the resource itself. If the condition of a resource is Normal, then its status is also Normal.

Tip: The **Service Classes** page in the GUI displays both the available capacity and the unavailable capacity for a service class. The Unavailable Capacity column shows the amount of storage that satisfies the requirements of the service class, but is unavailable because of the consolidated status of the storage system or its internal resources.

• If a service class allows provisioning only from certain capacity pools, any storage pool that is not in a specified capacity pool is eliminated as a candidate for provisioning.

Tip: You can see whether a service class restricts provisioning to one or more capacity pools from the **Service Classes page**. On the **Service Classes** page, show the Capacity Pools column, which is hidden by default.

• A storage pool must satisfy the requirements of every service class property in the following table to be a candidate for provisioning. Other properties of a service class are instructions for provisioning and do not affect whether the storage pool is a candidate for provisioning.

Table 219. Filtering candidates for provisioning based on service class properties	
Service class property:	Candidates for provisioning:
Storage tier	If specified in the service class, storage pools that are assigned to any storage tier within the specified range are candidates for provisioning.
RAID level	If a RAID level is specified, only storage pools that are configured at the specified RAID level are candidates for provisioning.
	Exception: The IBM XIV Storage System uses a nontraditional approach to data redundancy and does not support any of conventional RAID levels. IBM Spectrum Control always considers XIV storage pools to be configured at RAID 10.
	Restriction: Storage pools that support multiple RAID levels are not candidates for provisioning when a service class specifies a RAID level. Storage pools that support multiple RAID levels are not candidates even when the specified RAID level is one of the RAID levels that are supported by the storage pool. Storage pools that support multiple RAID levels are only candidates for provisioning when the service class specifies that any RAID level is allowed.

Table 219. Filtering candidates for provisioning based on service class properties (continued)	
Service class property:	Candidates for provisioning:
Virtualization	If virtualization is required by the service class, only storage pools on IBM storage virtualizers are candidates for provisioning. The IBM storage virtualizers are SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, FlashSystem 9100, FlashSystem V840, and FlashSystem V9000.
	If virtualization is not allowed by the service class, only storage pools that are not on IBM storage virtualizers are candidates for provisioning.
Encryption	If encryption is required by the service class, only storage pools on DS8000 storage systems that are encrypted and configured to use the encryption group that is specified in the service class are candidates for provisioning.
	If encryption is not allowed by the service class, only pools that are not encrypted are candidates for provisioning.
Thin provisioning	If thin provisioning is required by the service class, only storage pools that are configured for thin provisioning are candidates for provisioning. Thin provisioning is supported by the XIV, DS8000, and storage systems that run IBM Spectrum Virtualize. During provisioning, a thin-provisioned volume is created.
	If thin provisioning is not allowed by the service class, any storage pool that is not thin provisioned is a candidate for provisioning. In addition, any storage pool on DS8000 or a storage system that runs IBM Spectrum Virtualize is a candidate for provisioning, regardless of whether it is configured for thin provisioning. During provisioning, the volume that is created is not thin provisioned.
Compression	If compression is required by the service class, at least one I/O group of an IBM storage virtualizer must be enabled for compression. Only pools of the IBM storage virtualizer are candidates for provisioning.
Overallocation limit	If thin provisioning is required by the service class, IBM Spectrum Control identifies the storage pools that can thin provision volumes as candidates for provisioning. If an overallocation limit is specified, IBM Spectrum Control refines that list to include only the storage pools that have an overallocation percentage that does not exceed the overallocation limit.
	The overallocation percentage for a pool is determined by using the following measurements, as shown on the Pools page:
	 For an XIV storage pool, the Soft Capacity divided by the Capacity. For other storage system pools, the Provisioned Capacity.
Resource Tags	Resource tags further refine the candidates for provisioning. A service class can specify up to 3 resource tags. If resource tags are specified for the service class, only pools that have all the same tags are candidates for provisioning. If a pool is not tagged, any tags on the containing storage system also apply to the pool.

After the storage pools that can provide the service class are identified, IBM Spectrum Control calculates the amount of space that is available for provisioning new volumes of the service class. The calculated volume capacity is shown in the Available Capacity column on the **Service Classes** page.

The available-space calculation differs depending on whether the block-storage service class requires thin provisioning, as follows:

- If thin provisioning is not allowed, the available capacity is the sum of the following measurements, as shown on the **Pools** page:
 - For each XIV System pool, the Available Soft Capacity minus any Reserved Capacity
 - For other storage system pools, the Available Capacity minus any Reserved Capacity.
- If thin provisioning is required, the available capacity is the total of the following measurements, as shown on the **Pools** page, and calculations:
 - For each XIV System pool, the Available Soft Capacity minus any Reserved Capacity
 - For other storage system pools, the pool Capacity multiplied by the overallocation limit that is specified in the service class, minus the Provisioned Capacity. Any Reserved Capacity is subtracted from the result. For DS8000 storage pools, any space that is used by Track Space Efficient (TSE) volumes is subtracted from the calculated available capacity.

How the recommended placement of storage volumes is determined

When you request one or more new volumes by using the Provision Storage wizard, the candidates for provisioning are further refined and a provisioning task is created which shows the recommended placement of the new volumes. The same filtering that was used to determine the available capacity for the service class is used, and is refined further as follows:

- A provisioning request specifies the set of storage pools from which the storage can be provisioned. This set can be all the storage pools that are known to IBM Spectrum Control, or can be limited to the storage pools in a particular capacity pool. Only the storage pools in the specified set are candidates for provisioning. This set might be more restrictive than the set used to calculate the available capacity for the service class.
- Any storage pool that cannot provide the requested volume capacity is eliminated as a candidate for provisioning.
- If virtualization and VDisk mirroring are specified by the service class, then a VDisk copy of the volume must be made in a second storage pool on the same IBM storage virtualizer. If the storage pools on the IBM storage virtualizer cannot provide the requested capacity for both the volume and its mirrored copy, the pools are eliminated as candidates for provisioning.
- If an overallocation limit is specified in the service class, only the storage pools that will not exceed the limit after the new volume or volumes are provisioned are candidates for provisioning.

Remember: Storage pools that have an overallocation percentage that does not exceed the overallocation limit, are included in the available-space calculation. However, if the overallocation percentage would exceed the overallocation limit after the requested volume space is included, the pools are not candidates for provisioning.

After the storage pools that can satisfy the provisioning request are identified, IBM Spectrum Control finds the best location for the storage. The best location for the storage is based on the unused volume capacity in the pool and performance data. Preference is first given to storage pools and systems that already contain volumes for the selected server or hypervisor. Preference is given to systems that have available performance data.

File storage: Calculating available capacity and determining the placement of shares

For each file-storage service class, IBM Spectrum Control calculates the amount of available capacity for NAS file shares of that service class. When you request file shares of a particular service class, IBM Spectrum Control recommends the best location for the shares from the available capacity.

For each file-storage service class, IBM Spectrum Control calculates the amount of available capacity for NAS file shares of that service class. When you request file shares of a particular service class, IBM Spectrum Control recommends the best location for the shares from the available capacity.

The **Service Classes** page displays the available capacity for each file-storage service class. The available capacity is the amount of space that is available across all IBM Storwize V7000 Unified file systems that are known to IBM Spectrum Control and that can satisfy the requirements of the service class. The set of

file storage resources that are used to calculate the available capacity are also the candidates for provisioning by using the service class.

When you request a file share by using the Provision Storage pages, you specify the capacity that is required for the share, and you specify a service class. You can also specify a capacity pool from which shares must be provisioned. Based on the requested size of the share and whether a capacity pool is specified, the candidates for provisioning are further refined to include only the file storage resources that have enough available capacity and are in the specified capacity pool. From the remaining candidates for provisioning, IBM Spectrum Control creates a provisioning task that identifies the best location for the new share.

How available capacity for a file-storage service class is calculated

For each file-storage service class, IBM Spectrum Control calculates the available capacity for file shares of the service class. Only the file storage resources that are candidates for provisioning by using the service class are used to calculate the available capacity.

Identifying candidates for provisioning is a filtering process. File systems that could be candidates for provisioning if shared storage is allowed are eliminated as candidates for provisioning when dedicated storage is required. IBM Spectrum Control filters the set of available storage resources against each attribute of the service class until only the storage resources that can satisfy all of the requirements remain. These remaining storage resources are the candidates for provisioning.

The candidates for provisioning are identified by eliminating storage resources that cannot provide the service class, as follows:

- Only IBM Storwize V7000 Unified systems are supported for provisioning file shares. No other storage system type is a candidate for provisioning file shares.
- If a storage system has a consolidated status of Error or Unreachable, none of its internal NSDs or file systems are included in the available-space calculation.
- If an NSD or file system has a consolidated status of Error, it is not included in the available-space calculation.

Tip: The status of file systems is not displayed on the GUI. View the status on the NSD on which a file system resides to determine the file system status.

• If the storage volume or NSD on which a file system resides has a consolidate status of Error, the file system is not included in the available-space calculation.

Remember: Status is different from condition. Status is the status reported by a storage resource without considering the status of its internal resources. Condition is the combined status of the monitored, internal resources for the higher-level resource, and the status of the higher-level resource itself. The condition of a storage system is shown on the **Storage Systems** page, the condition of a volume is shown on the **Volumes** page, and the condition of an NSD is shown on the storage system details page. The condition might not have the same value as the status, which is used to eliminate file storage resources as candidates for provisioning.

Although the condition of a resource might not have the same value as its status, the condition is always at least as severe as the status. The condition reflects the most severe status of any internal resource or the status of the resource itself. If the condition of a resource is Normal, then its status is also Normal.

Tip: The **Service Classes** page displays both the available capacity and the unavailable capacity for a service class. The Unavailable Capacity column shows the amount of storage that satisfies the requirements of the service class, but is unavailable because of the consolidated status of the storage system or its internal resources.

• If a service class allows provisioning only from certain capacity pools, any storage resource that is not in a specified capacity pool is eliminated as a candidate for provisioning.

Tip: You can see whether a service class restricts provisioning to one or more capacity pools from the **Service Classes page**. On the **Service Classes** page, show the Capacity Pools column, which is hidden by default.

• A file system or NSD must satisfy the requirements of every service class property in the following table to be a candidate for provisioning. Other properties of a service class are instructions for provisioning and do not affect the storage placement determination.

Table 220. Filtering candidates for provisioning based on service class properties	
Service class attribute:	Candidates for provisioning:
Shared storage / Dedicated storage	If the service class specifies dedicated storage, the share can be created only on an unused NSD. File systems and other file storage resources are eliminated as candidates for provisioning.
Resource tags	Resource tags further refine the candidates for provisioning. A service class can specify up to 3 resource tags. If resource tags are specified for the service class, only file systems and NSDs that have all the same tags are candidates for provisioning. If a file system or NSD is not tagged, any tags on the containing storage system also apply to the internal resource.

After the file storage resources that can provide the service class are identified, IBM Spectrum Control calculates the amount of space that is available for providing new files shares of the service class.

The available-space calculation differs depending on whether the file-storage service class requires dedicated storage or allows shared storage, as follows:

- If the file-storage service class requires dedicated storage, the available capacity is the total available capacity on unused NSDs within each storage system. You can verify that an NSD is unused, and view its available capacity, from the details page of a storage system. To view information about the NSDs that are associated with a storage system, click **Network Shared Disks**. If the NSD is unused, the File Systems column shows 0. The Available Capacity column shows the amount of space on the NSD, which is the measurement that is used in the available-space calculation for the service class.
- If the file-storage service class allows shared storage, the file system in which the NAS share is provisioned can contain other NAS shares. In this case the available capacity is the total of the following measurements:
 - The available capacity on unused NSDs within each storage system.
 - The available capacity for each file system within each IBM Storwize V7000 Unified storage system. You can view the file system available capacity from the details page of a file storage system. To view information about the file systems that are in a storage system, click **File Systems**. The Available Capacity column shows the amount of available capacity on each file system, which is the measurement that is used in the available-capacity calculation for the service class.

How the recommended placement of file shares is determined

When you request a file share by using the Provision Storage pages, the candidates for provisioning are further refined and a provisioning task is created which shows the recommended placement of the new share. The same filtering that was used to determine the available capacity for the service class is used, and is refined further as follows:

- A provisioning request specifies the set of file storage resources from which the storage can be
 provisioned. This set can be all the storage resources that are known to IBM Spectrum Control, or can
 be limited to the storage resources in a particular capacity pool. Only the storage resources in the
 specified set are candidates for provisioning. This set might be more restrictive than the set used to
 calculate the available capacity for the service class.
- Any storage resource that cannot provide the requested capacity for the share is eliminated as a candidate for provisioning.

The shares are created on file systems or NSDs. After the file systems and NSDs that can satisfy the provisioning request are identified, IBM Spectrum Control identifies the best location for the storage. The best location for the storage is based on the available capacity on the file system or NSD.

Changing the default host definition for provisioned storage

Use the **setdscfg** command to change the default host definition for fibre-channel ports when provisioning storage from resources that run IBM Spectrum Virtualize.

You must have Administrator authority to use the **setdscfg** command.

When you provision storage from a storage system that runs IBM Spectrum Virtualize, which does not have an existing host definition for the fibre-channel ports, IBM Spectrum Control creates a new host definition during the provisioning task. The new host definition can access only one I/O group on the source device that you specified when you created the provisioning task. To allow the new host definition to access any I/O group on the source device, use the management application for the source device or use the IBM Spectrum Control setdscfg command.

To change the host definition to allow access to any I/O groups on the source device:

- 1. Start the IBM Spectrum Control CLI by issuing the **tpctool** command.
- 2. Issue the **setdscfg** command with the following values for the **-property** and **-context** parameters:
 - a) Specify Configuration.RestrictNewHostsToIogroup as the property value.
 - b) Specify DiskManager as the context value and specify false as the context parameter.

For example, to allow the new host definition to access any I/O groups on the source device of the provisioning task:

```
tpctool setdscfg -url localhost:9550 -user ***** -pwd *****
-property Configuration.RestrictNewHostsToIogroup -context DiskManager false
```

The value of Configuration.RestrictNewHostsToIogroup is set to false in the IBM Spectrum Control database. This removes the default restriction on the new host definition and allows access to any I/O groups on the provisioning source device.

Tip: To verify that the host definition restriction is changed, enter the following command:

```
tpctool getdscfg -url localhost:9550 -user ***** -pwd *****
-property Configuration.RestrictNewHostsToIogroup -context DiskManager
```

Provisioning storage with the IBM Spectrum Control GUI

You can provision volumes or shares to servers, hypervisors, and clusters in IBM Spectrum Control.

Discontinued support: Provisioning is no longer supported in IBM Spectrum Control. While the feature might still work in this release, it's recommended that you use another tool for your provisioning needs, when possible. For a complete list of discontinued features, see "Discontinued features in IBM Spectrum Control" on page xxx.

The prerequisites for provisioning storage are as follows:

- You must have the required license. "Actions that are available based on role" on page 26
- You must have Administrator privileges or permission to use one or more of the service classes that are created for block-storage or file-service provisioning.
- You must add and probe block-storage or file-storage resources, or both.

The following restrictions apply to volume provisioning:

- If the service class that is associated with the provisioning request requires multipathing or if automatic zoning is enabled, the fabrics and switches must be managed by IBM Spectrum Control. If the fabric is not managed by IBM Spectrum Control, the fabric-related configuration options are ignored when the provisioning task is executed.
- If you provision volumes to physical servers and servers that are virtual machines, you must create separate provisioning tasks using the GUI.
- If you provision multiple servers with storage volumes, the servers must all run on the same operating system. If the servers are virtual machines, the hypervisors that manage the servers must all run on the same operating system.

• If the provisioning request involves assigning multiple volumes to a resource, the volumes are sourced from pools in the same storage system.

NAS file shares can be provisioned only from Storwize V7000 Unified storage systems.

Specify the capacity that is required and the service class that you want to associate with the provisioning task. For shares, you must also specify at least one file access protocol. IBM Spectrum Control identifies the storage resources that can provide the storage capacity and that meet the criteria that are specified for the service class. From the set of storage resources that can provide the capacity and service class, IBM Spectrum Control identifies the best location for the storage and creates a provisioning task. Depending on the privileges that are associated with your role, you can save, execute, or schedule the provisioning task.

Provisioning volumes to servers

You can provision volumes to one or more servers. The GUI guides you through the steps of requesting capacity and specifying requirements by using a service class.

You can provision volumes to one or more servers. If a server is a virtual machine, the storage is provisioned to the hypervisor that is managing the server.

1. In the menu bar, go to Servers > Servers..

Tip: Alternatively, you can go to Advanced Analytics > Provisioning, and click Provision to Servers.

- 2. Select one or more servers, and then select **Actions** > **Provision Storage**.
- 3. Follow the instructions to assign the storage volumes.

The GUI enables you to create a provisioning task.

Related tasks

Provisioning storage with the IBM Spectrum Control GUI

You can provision volumes or shares to servers, hypervisors, and clusters in IBM Spectrum Control.

Provisioning volumes to hypervisors

You can provision volumes to one or more hypervisors. To provision volumes, the GUI guides you through the steps of requesting capacity and specifying requirements by using a service class.

You can provision volumes to one or more hypervisors.

1. Servers > Hypervisors

Tip: Alternatively, you can go to **Advanced Analytics** > **Provisioning**, and click **Provision to Hypervisors**.

- 2. Select one or more hypervisors, and then click **Actions** > **Provision Storage**.
- 3. Follow the instructions to assign the storage volumes.

The GUI creates a provisioning task.

Related tasks

Provisioning storage with the IBM Spectrum Control GUI

You can provision volumes or shares to servers, hypervisors, and clusters in IBM Spectrum Control.

Provisioning volumes to a server cluster

You can assign volumes to a cluster of servers with the GUI.

If the cluster is on servers that are virtual machines, the storage is assigned to the hypervisors that manage the servers. The volumes are assigned to all of the nodes in the cluster.

1. In the menu bar, go to **Servers** > **Servers**, and click **Clusters**.

Tip: Alternatively, you can go to **Advanced Analytics** > **Provisioning**, click **Provision to Servers**, and click **Clusters**.

- 2. Select a cluster, and then click **Actions** > **Provision Storage**.
- 3. Follow the instructions to assign the storage volumes.

The GUI enabled you to create a provisioning task.

Related tasks

Provisioning storage with the IBM Spectrum Control GUI

You can provision volumes or shares to servers, hypervisors, and clusters in IBM Spectrum Control.

Provisioning volumes to a hypervisor cluster

You can assign volumes to a cluster of hypervisors.

The volumes are assigned to all of the nodes in the hypervisor cluster.

1. In the menu bar, go to **Servers** > **Hypervisors**, and click **Clusters**.

Tip: Alternatively, you can go to **Advanced Analytics** > **Provisioning**, click **Provision to Hypervisors** and click **Clusters**.

- 2. Select a cluster, and then click **Actions** > **Provision Storage**.
- 3. Follow the instructions to assign the storage volumes.

The GUI created a provisioning task.

Related tasks

Provisioning storage with the IBM Spectrum Control GUI

You can provision volumes or shares to servers, hypervisors, and clusters in IBM Spectrum Control.

Provisioning shares to servers

You can provision a network-attached storage (NAS) file share to one or more servers. The GUI guides you through the steps of requesting capacity, specifying requirements by using a service class, and providing information about how to export the share.

You can provision a NAS file share to one or more servers. The share is made available to the selected servers, but is not automatically mounted from the selected servers.

1. In the menu bar, go to Servers > Servers..

Tip: Alternatively, you can go to Advanced Analytics > Provisioning, and click Provision to Servers.

- 2. Select one or more servers, and then select **Actions** > **Provision Storage**.
- 3. Follow the instructions to assign the share.

The GUI enabled you to create a provisioning task.

Related tasks

Provisioning storage with the IBM Spectrum Control GUI

You can provision volumes or shares to servers, hypervisors, and clusters in IBM Spectrum Control.

Provisioning shares to hypervisors

You can provision a network-attached storage (NAS) file share to one or more hypervisors. The GUI guides you through the steps of requesting capacity, specifying requirements by using a service class, and providing information about how to export the share.

You can provision a NAS file share to one or more hypervisors. The share is made available to the selected hypervisors, but is not automatically mounted from the selected hypervisors.

1. In the menu bar, go to **Servers** > **Hypervisors**.

Tip: Alternatively, you can go to **Advanced Analytics** > **Provisioning**, and click **Provision to Hypervisors**.

- 2. Select one or more hypervisors, and then select **Actions** > **Provision Storage**.
- 3. Follow the instructions to assign the share.

The GUI enabled you to create a provisioning task.

Related tasks

Provisioning storage with the IBM Spectrum Control GUI

You can provision volumes or shares to servers, hypervisors, and clusters in IBM Spectrum Control.

Provisioning shares to a server cluster

You can assign a network-attached storage (NAS) file share to a server cluster using the GUI.

You can provision a NAS file share to a server cluster. The share is made available to the servers in the cluster, but is not automatically mounted from the selected servers.

To provision a NAS file share to a server cluster, follow these steps:

1. In the menu bar, go to **Servers** > **Servers**, and click **Clusters**.

Tip: Alternatively, you can go to **Advanced Analytics** > **Provisioning**, click **Provision to Servers**, and click **Clusters**.

- 2. Select a cluster, and then select **Actions** > **Provision Storage**.
- 3. Follow the instructions to assign the share.

A provisioning task was created for a server cluster.

Related tasks

Provisioning storage with the IBM Spectrum Control GUI

You can provision volumes or shares to servers, hypervisors, and clusters in IBM Spectrum Control.

Provisioning shares to a hypervisor cluster

You can assign a network-attached storage (NAS) file share to a hypervisor cluster.

You can provision a NAS file share to a hypervisor cluster. The share is made available to the hypervisors in the cluster, but is not automatically mounted from the selected hypervisors.

To provision a NAS file share to a hypervisor cluster, follow these steps:

1. In the menu bar, go to **Servers** > **Hypervisors**, and click **Clusters**.

Tip: Alternatively, you can go to **Advanced Analytics** > **Provisioning**, click **Provision to Hypervisors** and click **Clusters**.

- 2. Select a cluster, and then select **Actions** > **Provision Storage**.
- 3. Follow the instructions to assign the share.

A provisioning task was created for a hypervisor cluster.

Related tasks

Provisioning storage with the IBM Spectrum Control GUI

You can provision volumes or shares to servers, hypervisors, and clusters in IBM Spectrum Control.

Provisioning with zone control

When you provision storage, you can set a zoning policy to enable automatic zoning. When automatic zoning is enabled, IBM Spectrum Control can create zones during storage provisioning to connect a server to a storage system.

Discontinued support: Provisioning is no longer supported in IBM Spectrum Control. While the feature might still work in this release, it's recommended that you use another tool for your provisioning needs, when possible. For a complete list of discontinued features, see "Discontinued features in IBM Spectrum Control" on page xxx.

When IBM Spectrum Control creates a provisioning task, it identifies the best location for the new storage that satisfies the requirements of the service class. If automatic zoning is enabled, then, during provisioning, existing zones are used if the server already has connectivity to the storage system. Otherwise, one or more zones are created between a host initiator port and a controller, node, or module port.

IBM Spectrum Control provides the support for zone control and management of Brocade switches and fabrics that are managed by using Brocade Network Advisor, and Cisco switches and fabrics.

The Cisco support includes:

• Standard Cisco zone control capabilities

· Cisco Enhanced Zoning features.

To set or modify the zoning policy, complete the following steps:

- 1. In the menu bar, go to Advanced Analytics > Provisioning, and click Set Zoning Policy.
- 2. Enable or disable automatic zoning:
 - Enable automatic zoning to allow IBM Spectrum Control to create zones during storage provisioning.
 The naming convention for new zones is host_storage-system_suffix-number. When you enable automatic zoning, you can specify the following options:

Zone name prefix

If you specify a zone name prefix, all zones that are automatically created by IBM Spectrum Control are prefixed with a string you specify. This prefix can help you identify which zones were automatically created, and which were created manually.

Make changes to the active zone set

Specify whether changes are made to the active zone set or to a new inactive zone set. If this check box is selected, one or more new zones are added to the active zone set. If the check box is cleared, a new inactive zone set is created and the new zone or zones are added to this new inactive zone set. The new inactive zone set will contain only the new zones.

- Disable automatic zoning if you want IBM Spectrum Control to use only existing zones. If you disable automatic zoning, no changes will be made to the zoning for the fabric.
- 3. Click Save.

Optimizing storage tiering

To optimize the placement of volumes on storage tiers, analyze the tiering of volumes in your storage environment.

Discontinued support: Storage optimization is no longer supported in IBM Spectrum Control. While the feature might still work in this release, it's recommended that you use another tool for your optimization and tiering needs, when possible. For a complete list of discontinued features, see "Discontinued features in IBM Spectrum Control" on page xxx.

Before you analyze tiering, complete the following tasks:

- · Probe the storage virtualizers
- · Collect performance data
- Set the tier level of the storage pools that you analyze and that you select as target pools

Analyzing tiering 1 Collect data \hookrightarrow Probe the storage virtualizers to collect information about the logical devices, such as the volumes and virtual disks, that are associated with the storage virtualizer. 2 Collect performance data 1010 Collect performance data to gauge the average workload activity of the volumes that you want to analyze. 3 Assign tier levels Assign tier levels to the pools that contain the volumes that you want to analyze and to the target pools. 4 Choose the source storage resources Select the resources that you want to analyze: → Servers ⇔ Storage virtualizers → Storage pools → Volumes **5 Choose the target pools** You select the target pools to place volumes that meet the tiering criteria that you specify. 6 Choose the performance data Choose the analysis period, the collection days, and time period that you want to use to analyze the tiering of the volumes. 7 Specify the tiering criteria Specify the tiering criteria and the pool saturation rates for the target pools.

- 1. "Collecting data" on page 138
- 2. "Collecting performance data by using IBM Spectrum Control performance monitors" on page 144
- 3. "Setting the tier level of storage pools" on page 441

- 4. "Analyzing tiering by servers" on page 442
- 5. "Analyzing tiering by hypervisors" on page 442
- 6. "Analyzing tiering by storage virtualizers" on page 442
- 7. "Analyzing tiering by storage pools" on page 443
- 8. "Analyzing tiering by volumes" on page 443
- 9. "Tiering volumes by I/O density and I/O rate" on page 443
- 10. "Tiering volumes by I/O density and I/O rate" on page 443

The purpose of analyzing tiering is to tier volumes based on the I/O rate or I/O density thresholds that you set. Depending on the conditions that are set, recommendations are generated.

Storage costs can be reduced, for example, by moving volumes with low workloads from higher and more expensive tiers to lower and less expensive tiers. Storage performance can be improved and used more efficiently by moving volumes with heavy workloads to the tiers that meet the workload requirements of the volumes. When you implement the recommendations, you ensure that the volumes are placed on the tiers that best match the workload requirements of the volumes.

You can optimize the placement of volumes on storage systems that run IBM Spectrum Virtualize.

Restriction: You cannot optimize the placement of volumes in data reduction pools.

Learn more: For information about the components that can be used with IBM Spectrum Control Version 5.2 (or later), see

IBM Spectrum Control interoperability matrix at http://www.ibm.com/support/docview.wss?uid=swg21386446.

Volumes can be moved to tiered storage pools on the same storage virtualizer, but volumes cannot be moved from one storage virtualizer pool to another storage virtualizer pool.

Optimizing volumes with service classes: Recommendations are not generated to move volumes that are assigned a service class to a destination pool that does not meet the requirements of the service class.

Related tasks

Analyzing and re-tiering volumes in pools on tier 1

In this tutorial task, you analyze the current placement of volumes in pools on tier 1 so that you can move volumes with low workloads from tier 1 to tier 2 or tier 3 pools.

"Tutorial: Collocating volumes" on page 90

In this tutorial task, you want to minimize the exposure of servers to multiple back-end storage systems by collocating volumes that are assigned to the same hypervisor or server. You can enforce the collocation of volumes when you enter the criteria for analyzing tiering and balancing pools.

Investigating the capacity of tiered storage

Review the capacity that is assigned to tiered storage to see whether your system needs more storage space or whether you can re-tier storage to satisfy the capacity demands for tiered storage. View the space usage trends at the current growth rate for each tier and the projected depletion date of storage space.

Add the pools in your storage environment to the tiers that best meet the performance requirements of your storage data.

You can view the total capacity by tier and the capacity that is consumed by the pools that are not tiered. You can also view information about recent storage growth for each tier, including the weekly and average storage growth, and the date when the storage space of the tier is expected to run out.

- 1. Click **Groups** > **Tiers**.
- 2. Review the information that is shown about the distribution of capacity across each tier.
- 3. Optional: To place pools on lower or higher tiers, complete these steps:
 - a) Click the number in the **Pools** column.
 - b) Select the pools that you want to re-tier.

- c) Right-click the pools, click **Set Tier**, and click the tier level.
- 4. Optional: To assign untiered pools to tiers, click **Untiered Pools**.

If the storage space for tier 1 pools is running out, you can down-tier the volumes in the tier-1 pools to lower tiers. Alternatively, you can add more storage to meet the requirements of your tier-1 storage data.

Related tasks

Setting the tier level of storage pools

To analyze tiering and to balance pools, you must set the tier level of storage pools.

Analyzing tiering by storage virtualizers

Optimize tiering by analyzing the tiering of volumes on storage virtualizers.

Analyzing tiering by volumes

Optimize tiering by analyzing the tiering of volumes on storage virtualizers.

Tiering volumes by I/O density and I/O rate

Set thresholds for tiering volumes by I/O rate or I/O density. To gain better optimization results, you can also customize the analysis of the workload activity of the pools.

Renaming tiers

Change the default names of the tiers so that they match the names of the tiers in your storage environment.

Setting the tier level of storage pools

To analyze tiering and to balance pools, you must set the tier level of storage pools.

Before you analyze tiering or balance pools, you must set the tier level of pools.

To analyze tiering, you must set the tier level of the source pools and the target pools. The source pools are the pools that you select for analysis. The target pools are the pools that are used to relocate volumes to lower or higher tiers of storage.

Tip: You can also set or modify the tier level of pools on the **General** tab of the **View Properties** page for the storage pool.

Child pools have the same tier level as the parent pool. So if you change the tier level of the parent pool or of the child pool, all of the pools in the parent-child relationship are set to the same tier level.

You can assign tier levels to storage systems that are not storage virtualizers. However, you cannot balance the pools or analyze the tiering of pools other than storage virtualizer pools.

- 1. From the **Storage** menu, click **Pools**.
- 2. On the **Pools** page, select one or more storage pools.
- 3. Click **Set** > **Tier**, and then select a tier level.

To remove the tier level, click None.

The tier level that you selected for the storage pools is shown in the **Tier** column.

Related tasks

Investigating the capacity of tiered storage

Review the capacity that is assigned to tiered storage to see whether your system needs more storage space or whether you can re-tier storage to satisfy the capacity demands for tiered storage. View the space usage trends at the current growth rate for each tier and the projected depletion date of storage space.

Renaming tiers

Change the default names of the tiers so that they match the names of the tiers in your storage environment.

Renaming tiers

Change the default names of the tiers so that they match the names of the tiers in your storage environment.

- 1. From the Groups menu, click Tiers.
- 2. On the Tiers page, click Rename Tiers.
- 3. On the **Rename Tiers** property notebook, click **Edit**.
- 4. Set the display names for your tiers and click **Save**.

Related tasks

Investigating the capacity of tiered storage

Review the capacity that is assigned to tiered storage to see whether your system needs more storage space or whether you can re-tier storage to satisfy the capacity demands for tiered storage. View the space usage trends at the current growth rate for each tier and the projected depletion date of storage space.

Setting the tier level of storage pools

To analyze tiering and to balance pools, you must set the tier level of storage pools.

Analyzing tiering by servers

Optimize tiering by analyzing the tiering of volumes on storage virtualizers that are connected to servers.

The servers that you select must be connected to one or more storage virtualizers.

- 1. From the **Servers** menu, click **Servers**.
- 2. Right-click one or more servers, and then click **Analyze Tiering**. Alternatively, you can click **Actions** > **Analyze Tiering**.

Analyzing tiering by hypervisors

Optimize tiering by analyzing the tiering of volumes on storage virtualizers that are connected to hypervisors.

The hypervisors that you select must be connected to one or more storage virtualizers.

- 1. From the **Servers** menu, click **Hypervisors**.
- 2. Right-click one or more hypervisors, and then click **Analyze Tiering**. Alternatively, you can click **Actions** > **Analyze Tiering**.

Analyzing tiering by storage virtualizers

Optimize tiering by analyzing the tiering of volumes on storage virtualizers.

- 1. From the Storage menu, click Storage Systems.
- 2. Right-click one or more storage virtualizers, and then click **Analyze Tiering**. Alternatively, you can click **Actions** > **Analyze Tiering**.

Related tasks

Investigating the capacity of tiered storage

Review the capacity that is assigned to tiered storage to see whether your system needs more storage space or whether you can re-tier storage to satisfy the capacity demands for tiered storage. View the

space usage trends at the current growth rate for each tier and the projected depletion date of storage space.

Analyzing tiering by storage pools

Optimize tiering by analyzing the tiering of volumes in storage pools.

- 1. From the **Storage** > **Pools** menu, click **Pools**.
- 2. Right-click one or more storage pools, and then click **Analyze Tiering**. Alternatively, you can click **Actions** > **Analyze Tiering**.

Analyzing tiering by volumes

Optimize tiering by analyzing the tiering of volumes on storage virtualizers.

- 1. From the **Storage** menu, click **Volumes**.
- 2. Right-click one or more volumes, and then click **Analyze Tiering**. Alternatively, you can click **Actions** > **Analyze Tiering**.

Related tasks

Investigating the capacity of tiered storage

Review the capacity that is assigned to tiered storage to see whether your system needs more storage space or whether you can re-tier storage to satisfy the capacity demands for tiered storage. View the space usage trends at the current growth rate for each tier and the projected depletion date of storage space.

Tiering volumes by I/O density and I/O rate

Set thresholds for tiering volumes by I/O rate or I/O density. To gain better optimization results, you can also customize the analysis of the workload activity of the pools.

- 1. Select the resources that you want to analyze.
 - The source storage pools that are related to the resources that you selected are analyzed to determine whether they meet the workload requirements of the volumes. If the workload requirements of the volume in its current tier are not met, the volume is a candidate for relocation.
- 2. Select the target storage pools.
- 3. Include or exclude volumes in mirrored volume relationships from the analysis.
 - **Mirrored volumes:** If you want to optimize mirrored volumes, select one or more target pools that do not contain either the primary volume or the secondary volume. Primary volumes and secondary volumes cannot be placed in the same destination pool.
- 4. Select the period that you want to use to analyze the performance data.
 - You can choose the number of days, days of the week, and hours of the day for the analysis to occur.
- 5. Specify whether volumes that are in the same source storage pool and that are assigned to the same server or hypervisor must be placed in the same destination storage pool.
 - Recommendations are only generated to move one or more of the related volumes when all of the related volumes can be placed in the same destination storage pool.
 - **Multiple host connections to the same hypervisor or server:** If the volumes in the source pool that are assigned to the same hypervisor or server are assigned to different host connections, the collocation of the volumes is affected. In such cases, if volumes require optimization, the volumes that are assigned to the same host connection are kept together. To view information about the host connection for the volume, right-click the volume, select **View Details**, and then click the **Host Connections** tab
- 6. Specify the type of threshold and the threshold values for moving volumes to higher and lower tiers.
- 7. Click Analyze.

When you complete entering the information to analyze the tiering of the volumes, recommendations for tiering are shown on the Tiering Analysis page. You can create a schedule to run the analysis that is based on the analysis criteria that you entered.

Related tasks

Investigating the capacity of tiered storage

Review the capacity that is assigned to tiered storage to see whether your system needs more storage space or whether you can re-tier storage to satisfy the capacity demands for tiered storage. View the space usage trends at the current growth rate for each tier and the projected depletion date of storage space.

Modifying the criteria for analyzing tiering

When a task is created after you run the tiering analysis, you can modify the criteria for tiering the volumes.

You must run the tiering analysis to create a tiering analysis task.

The criteria that you modify are used to analyze the source volumes and target volumes that were selected when you created the tiering analysis.

- 1. From the Home menu, click Tasks.
- 2. Right-click a tiering analysis task, and click **Edit**.
- 3. Modify the tiering thresholds, and then click **Analyze**.

The source volumes are analyzed to determine whether one or more of the volumes require re-tiering based on the criteria that you specified. If one or more of the volumes require re-tiering, recommendations are generated to move the volumes to the destination pools that meet the criteria that you specified.

Optimizing storage pools

Analyze the activity of pools and resolve performance hot spots by redistributing volumes across each storage tier.

Discontinued support: Storage optimization is no longer supported in IBM Spectrum Control. While the feature might still work in this release, it's recommended that you use another tool for your optimization needs, when possible. For a complete list of discontinued features, see "Discontinued features in IBM Spectrum Control" on page xxx.

Before you balance pools, complete the following tasks:

- Probe the storage virtualizer. You must probe the storage virtualizer to collect information about the logical devices, such as volumes and virtual disks, that are associated with the storage virtualizer.
- Collect performance data. To gauge the average activity level of the volumes and pools, you must collect sufficient performance data for the storage virtualizer.
- Assign a tier level to the pools that you want to balance.

You can balance the workload of volumes in pools on the same tier on storage systems that run IBM Spectrum Virtualize.

Restriction: You cannot balance the workload of volumes in data reduction pools.

Learn more: For information about the components that can be used with IBM Spectrum Control Version 5.2 (or later), see IBM Spectrum Control interoperability matrix at http://www.ibm.com/support/ docview.wss?uid=swg21386446.

To balance pools, you must select two or more pools on the same tier level and on the same storage virtualizer. For example, you select three tier 1 pools, and four tier 3 pools. The pools on tier 1 are analyzed and balanced, and the pools on tier 3 are analyzed and balanced.

Optimizing volumes with service classes: Recommendations are not generated to move volumes that are assigned a service class to a destination pool that does not meet the requirements of the service class.

Related information

Setting the tier level of storage pools

Balancing pools

You balance pools to distribute the workload of volumes across pools on the same tier and on the same storage virtualizer. To improve the performance of the pools, recommendations are generated to move volumes to other pools on the same tier and on the same storage virtualizer.

- 1. From the **Storage** menu, click **Pools**.
- 2. Right-click two or more pools on the same tier level and on the same storage virtualizer, and then click **Balance Pools**.

Tip: If you want to analyze pools that contain volumes in mirrored volume relationships, select one or more pools that do not contain either the primary volume or the secondary volume. The primary volume and the secondary volume cannot be placed in the same destination pool.

- 3. Choose whether to analyze volumes in mirrored volume relationships.

 These options are only available if you select one or more volumes with mirrored volumes.
- 4. Select the period that you want to use to analyze the performance data.
- 5. Specify the placement of optimized volumes in capacity pools.

 These options are only available if you select one or more volumes in capacity pools.
- 6. Specify whether volumes that are in the same source storage pool and that are assigned to the same server or hypervisor must be placed in the same destination storage pool.
- 7. Click Analyze.

When you finish entering the information for balancing the pools, a task is created that shows the recommendations on the Balance Analysis page. You can implement the recommended changes, or you can create a schedule to implement the changes later.

Analyzing volumes with service classes: Recommendations are not generated to move volumes that are assigned a service class unless the destination pool can meet the requirements of the service class.

Modifying the criteria for balancing pools

When a task is created after you run the analysis for balancing pools, you can modify the criteria that determine which pools are balanced.

To create a balance analysis task, you must run the analysis for balancing the pools.

The criteria that you modify are used to analyze the storage pools that you selected when you entered the criteria for balancing the pools.

- 1. From the Home menu, click Tasks.
- 2. Right-click a balance analysis task, and click Edit.
- 3. Make your changes, and then click **Analyze**.

Based on the criteria that you specified, the pools are analyzed to determine whether the pools require balancing. If one or more pools require balancing, recommendations are generated to redistribute the volumes to pools on the same tier of storage.

Criteria for identifying the pools that require balancing

To identify the pools that require balancing, performance data is collected to determine the activity of pools on the same tier. The performance capability of pools is preserved by setting an activity limit for each tier of storage.

To determine the activity of pools on the same tier and on the same back-end storage system, the current activity level of the pool is calculated.

Tip: On the **Pools** page, the activity level of the pool is shown in the **Activity** column.

The following formulas are used to calculate the activity level of pools.

Table 221. Formulas for calculating the activity level of pools	
Value calculated	Formula that is used in the calculation
Activity level of the pool	<pre>[Read I/O Rate × (1- Read I/O Cache Hit %) + Write I/O Rate] ÷ Total Pool Capacity</pre>
Activity level of the pool on XIV systems	(Total I/O Rate÷Total Capacity)

Reclaiming storage

Use the recommendations to reclaim capacity before you plan new capacity purchases.

To determine which volumes can be reclaimed, capacity and performance data must be available for the reclamation analysis. To run the reclamation analysis, capacity data is collected for the previous day, and a daily aggregation of the performance data is collected for the previous 14 days.

Tip: Depending on the I/O activity of the volumes in your storage environment, you can shorten or lengthen the default period of 14 days.

Storage systems that support reclamation: To find out, go to BM Spectrum Control interoperability matrix, and then click the release number in the **Storage** column.

Got an alert that your storage system is running out of capacity? Or, did you click **Storage** > **Pools**, had a look at the values in the **Zero Capacity** column and saw that you were running out of capacity?

Instead of purchasing more capacity, click **Advanced Analytics** > **Reclamation** to see how much capacity you can reclaim.

The reclamation analysis is run daily and if sufficient data is collected recommendations are generated to reclaim the volumes that meet either one of the following criteria:

- The volume isn't assigned to a server.
- I/O activity was not detected in the data that was collected for the volume.

For example, the reclamation analysis is run, and it is determined that a volume isn't assigned to a server. The volume is identified as reclaimable even if I/O activity is detected for the volume. Alternatively, the reclamation analysis is run, and no I/O activity is detected for the volume. The volume is identified as reclaimable even if the volume is assigned to a server.

The reclamation analysis also detects if volumes are replica volumes, VDisk mirrored volumes, or FlashCopy volumes. The same criteria are used to identify whether the volumes are reclaimable, but the process varies depending on the type of copy volume:

Replica volumes

The source volumes of volumes with replicas, such as volumes that use Metro Mirror or Global Mirror copy services, are analyzed to determine whether the volumes are reclaimable.

VDisk mirrored volumes

Both copies of the VDisk mirrored volume are analyzed to determine whether the volumes are reclaimable.

FlashCopy volumes

The target volumes of volumes in FlashCopy relationships are analyzed if the source volumes are identified as reclaimable.

After the analysis is run, you see the total capacity that can be reclaimed, which is broken down by tier, and by volume. (The reclaimable capacity for volumes in storage systems that use data reduction technologies isn't shown. See the restrictions below.) In the table, you get a list of the volumes, the volume's capacity, and the information that you need to decide whether you want to decommission the volumes.

Tip: See volumes that are identified as reclaimable, such as volumes that are used to back up data or volumes that you recently assigned to a new application? If you don't want to include these volumes in the reclamation analysis, right-click the volumes and click **Exclude from Analysis**.

Restriction:

The following restrictions apply:

- For volumes on IBM Spectrum Accelerate and CKD volumes on DS8000, the volumes are identified as
 reclaimable based on I/O activity, because information about the assignment of volumes to servers is
 not available.
- Volumes in storage systems that use data reduction technologies are identified as reclaimable, but the
 actual physical capacity of the volumes can't be determined. For example, when data is deduplicated,
 multiple volumes can share identical blocks of data so we don't know the actual capacity of each
 volume. Because we can't determine the capacity of the individual volumes, the total reclaimable
 capacity for each volume is displayed as zero. This applies to:
 - Volumes in IBM storage systems that support data reduction pools or volumes in storage systems that run IBM Spectrum Virtualize.
 - Volumes in IBM FlashSystem A9000 storage systems.

Private volumes in Dell EMC storage systems are excluded from the reclamation analysis.

Viewing storage reclamation

Use the storage reclamation analysis to see information about the reclaimable capacity in your data center. You can see the savings that can be made by reclaiming capacity for tiered and non-tiered storage and view a list of the reclaimable volumes.

Before you view recommendations for reclaiming storage, complete the following tasks:

- Add the storage systems for monitoring and schedule a storage systems probe.
- If your storage environment is configured for storage tiers, ensure that the amount of available capacity on tier one is minimal.

By default, the period for analyzing performance data, which is used for generating reclamation recommendations, is set to 14 days. For a more thorough analysis, you can change this period.

View information about volumes that are analyzed and recommended for reclamation.

- 1. From the Advanced Analytics menu, click Reclamation.
- 2. On the **View by Reclaimable Capacity** page, view the charts and tables for information about reclaimable storage space, volumes that are recommended for reclamation, and volumes that are excluded from analysis.

On the reclamation donut chart, you can view the amount of storage space that is used and the amount of storage space that can be reclaimed. You can also see an estimate of the storage space that can be saved when the volumes that are listed in the table are reclaimed.

On the reclamation by tier bar charts, view the amount of storage space that can be reclaimed on each tier of storage that is defined in your data center.

On the **Recommendations** tab, view the volumes that are identified as potential candidates for reclamation.

On the **Excluded** tab, view the volumes that are excluded from the analysis for reclamation recommendations.

Viewing storage reclamation by storage system

Use the reclamation view of storage systems for detailed information about the reclaimable capacity in your data center. You can see the savings that can be made by reclaiming capacity for tiered and non-tiered storage and view information about the storage systems that contain the reclaimable volumes.

Before you view recommendations for reclaiming storage, complete the following tasks:

- Add the storage systems for monitoring and schedule a storage systems probe.
- If your storage environment is configured for storage tiers, ensure that the amount of unused capacity on tier one is minimal.

By default, the period for analyzing performance data, which is used for generating reclamation recommendations, is set to 14 days. For a more thorough analysis, you can change this period.

View reclamation information about storage systems that contain volumes that are analyzed and recommended for reclamation.

- 1. From the **Advanced Analytics** menu, click **Reclamation**.
- 2. On the **View by Storage Systems** page, view the charts and table for information about reclaimable storage space and for information about the storage systems that contain the reclaimable volumes.

On the reclamation donut chart, you can view the amount of capacity that can be reclaimed and the amount of capacity that can't be reclaimed. If you defined storage tiers, you can view the amount of storage that can be saved for each tier. You can also see an estimate of the storage space that can be saved when the volumes of the storage systems that are listed in the table are reclaimed.

On the reclamation bar chart, view the amount of storage space that can be reclaimed on each storage system in your data center.

In the table, view space, capacity, and reclamation information about each storage system.

Excluding volumes from reclamation analysis

You can select volumes to not include them in the analysis for reclamation recommendations. The volumes are excluded from the reclamation charts.

On the **View by Reclaimable Capacity** page, you can select a volume or multiple volumes for exclusion from the analysis for reclamation recommendations.

- 1. From the Advanced Analytics menu, click Reclamation.
- 2. On the **View by Reclaimable Capacity** page, to exclude a volume from the analysis, right-click the volume on the **Recommendations** table, and select **Exclude from Analysis**.

The volume is removed from the **Recommendations** table and added to the **Excluded** table, and the charts are refreshed.

Tip: To include a volume in the analysis, on the **Excluded** table, right-click the volume and select **Include in Analysis**. The volume is removed from the **Excluded** table and added to the **Recommendations** table, and the charts are refreshed.

Transforming and migrating volumes

You can transform fully allocated volumes to compressed or thin-provisioned volumes, or transform compressed or thin-provisioned volumes to fully allocated volumes. You can move volumes to other pools or to pools that are enabled for Easy Tier.

Discontinued support: Storage optimization is no longer supported in IBM Spectrum Control. While this feature might still work in this release, it's recommended that you use another tool for your volume transformation needs, when possible. For a complete list of discontinued features, see "Discontinued features in IBM Spectrum Control" on page xxx.

Before you transform storage, you must probe the storage virtualizers. The probe collects information about the logical devices, such as volumes and virtual disks, that are associated with the storage virtualizer.

You can convert or move volumes in pools on storage systems that run IBM Spectrum Virtualize.

For volumes in mirrored volume relationships, you can convert the primary volume, the secondary volume, or both volumes.

Restrictions: The restrictions for converting or moving volumes are as follows:

- You cannot convert or move volumes in data reduction pools.
- To move volumes, the destination pool must be on the same storage virtualizer. You cannot move a volume from a pool on one storage virtualizer to a pool on another storage virtualizer.
- To convert or move volumes, image mode volumes must be converted to managed mode volumes.
- To convert fully allocated volumes to compressed volumes, you must have the IBM Real-time Compression license.
- 1. In the menu bar in the web-based GUI, go to **Storage** > **Volumes**.
- 2. Click one or more storage volumes, and then select **Transform Storage** from the context-sensitive menu.

Alternatively, you can click **Actions** > **Transform Storage**.

When you complete entering criteria for transforming storage, the Transform Plan page is opened. A task is created which you can run to implement the recommendations.

Related tasks

"Modifying the period for analyzing performance data" on page 449

To change the default period for analyzing performance data in storage virtualizer pools, issue the **setdscfg** command. You can extend or shorten the default period for a more accurate analysis.

Modifying the period for analyzing performance data

To change the default period for analyzing performance data in storage virtualizer pools, issue the **setdscfg** command. You can extend or shorten the default period for a more accurate analysis.

You must have Administrator authority to use the **setdscfg** command.

By default, the period for analyzing performance data is set to 14 days. When you change the number of days, the period for analyzing performance data is modified for the advanced analytics about storage, that is, tiering, balancing, transforming, and reclamation.

When you transform storage, the performance data that is collected about the pools on storage virtualizers is analyzed. The period that is set for analyzing performance data is used to calculate the average workload activity of the volumes that are selected for analysis. To provide a more accurate analysis of the workload activity of the volumes in your storage environment, you can extend or shorten the period.

For reclamation, the performance data that is collected for the analysis period about all block storage pools is analyzed, regardless of whether the storage is virtualized.

To change the analysis period, complete the following steps:

- 1. To start the IBM Spectrum Control CLI, issue the **tpctool** command.
- 2. Issue the **setdscfg** command, and then set the following values by using the **-property** and the **-context** parameters:
 - a) Specify tiering_pm_days as the property value.
 - b) Specify *tiering* as the context value, and then specify the number of days for the analysis period. For example, to change the analysis period from the default value of 14 days to 28 days, issue the following command:

```
tpctool setdscfg -url localhost:9550 -user ***** -pwd *****
-property tiering_pm_days -context tiering 28
```

The default value is changed to the number of days that you specified, for example 28.

Tip: To verify that the analysis period is changed, you enter the following command:

```
tpctool getdscfg -url localhost:9550 -user ***** -pwd ***** -context tiering
```

Modifying the properties of resources

Add or change the properties for individual resources or for multiple resources. You can use the properties to filter or sort the resources in the GUI or in an external application if the data is shared or exported.

To add or modify the properties of resources, you must have Administrator privileges.

For servers, storage systems, hypervisors, switches, and fabrics, you can modify the properties of a single resource or multiple resources. You can modify the properties of a single pool, but not the properties of multiple pools.

Modifying the properties of multiple resources

- 1. From the menu, click the type of resource that you want to edit. For example, if you want to edit the properties for servers, click **Servers** > **Servers**.
- 2. To select the resources, click Ctrl + click or Shift + click.
- 3. Right-click the selected resources and click Edit Properties.

In the Edit Properties property notebook, the properties can have the following values:

- If the property is blank, it means that the resources that were selected have different values assigned to that property. For example, the selected servers have different locations.
- If the selected resources have the same value for a property, the value is shown.
- 4. Modify the properties and click Save.

Modifying the properties of a single resource

- 1. From the menu, click the type of resource that you want to edit. For example, if you want to edit the properties for servers, click **Servers** > **Servers**.
- 2. Right-click the resource that you want to modify and click **View Properties**.
- 3. In the property notebook, click **Edit**.
- 4. Modify the properties and click Save.

Properties of resources that can be modified

You can modify the properties, such as the name or location, of the resources that are monitored by IBM Spectrum Control.

The following table lists the properties that you can modify for each type of resource:

Resource	Properties that can be modified
Storage systems	Name (label)
	Location
	Custom tags (3)

Resource	Properties that can be modified
Pools	Tier Back-end storage system type Back-end storage RAID level Back-end storage disk type Back-end storage disks Tip: You can edit the back-end values if the backend storage system for the related storage virtualizer was not probed and the values for the pool are unknown. If you manually enter values for these properties, IBM Spectrum Control uses the values to help calculate the approximate read I/O capability of the pool.
Tip for editing the properties of multiple servers: If you select agentless servers and servers that are monitored by a Storage Resource agent, only the location and custom tag properties are shown on the Edit Properties property notebook. For servers that are monitored by a Storage Resource agent, the OS type property is automatically set when the server is probed and cannot be changed.	Servers with Storage Resource agents Location Custom tags (3) Agent trace Trace level Number of trace files Trace file size Run scripts on agent Servers without Storage Resource agents (agentless) You can modify the following properties for agentless servers: Name (label) OS type IP address Location Custom tags (3)
Hypervisors	Location Custom tags (3)
Switches	Name (label)LocationCustom tags (3)
Fabrics	Name (label) Location Custom tags (3)

Opening the management GUI for a resource

Resources such as storage systems and switches can have their own management GUIs or element managers. In IBM Spectrum Control, you can open the start page in these GUIs from resource list pages, such as the Storage Systems page and Switches page.

The action for opening a management GUI is available for storage systems and switches under the following conditions:

- The resource is monitored by IBM Spectrum Control.
- The resource has a management GUI that is web-based.
- In the menu bar in the GUI, go to a type of top-level resource.
 For example, if you want open the GUI for a block storage system, go to Storage > Block Storage
 Systems. If you want open the GUI for a switch, go to Network > Switches.
- 2. Right-click the resource and select the action to open the GUI. For storage systems, select **Launch Storage System GUI**. For a switch, select **Open Switch GUI**.

The start page for the management GUI is opened in a separate web browser window.

Opening the DS Storage Manager GUI for a DS8000

If the firmware level for a DS8000 is earlier than R6.2, you must complete extra configuration steps before you can open its management GUI. To determine the firmware version of a DS8000 and enable the ability to open its management GUI, complete the following steps:

- 1. Start IBM Spectrum Control.
- 2. In the menu bar, go to **Storage** > **Block Storage Systems**.
- 3. Right-click the DS8000 and select View Properties.
- 4. Check the version in the **Firmware** field. If the firmware version is earlier than R6.2, continue to the next step. If the version is R6.2 or later, you do not need to complete this procedure.
- 5. Download and start the PuTTY utility. PuTTY is a free implementation of Telnet and SSH for Windows and UNIX. You can download PuTTY from http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html.
- 6. On the **Session** page in the PuTTY tool, complete the following actions:
 - a) In the **Host Name (or IP address field)**, enter the IP address of the Hardware Management Console (HMC) for the DS8000. You can determine this IP address by viewing the properties of the DS8000 in the IBM Spectrum Control GUI.
 - b) Select SSH for a Connection type.
 - c) Click Open.
- 7. Enter the user name and password for logging in to the HMC.
- 8. Run the following command: touch /opt/esshmc/data/config/ntvbrsen
- 9. In the GUI, right-click the DS8000 and select Launch Storage System GUI.

Exporting information to a file

You can save information about resources, tasks, or alerts to a PDF, CSV, or HTML file. Information that you can export from the GUI includes all the values that are being shown in the columns for a list of resources, tasks, or alerts.

The information that you export to a file is organized according to the sorting, filtering, and column order that is defined for a list. Before you export information, complete the following tasks to configure a list:

- Change the order of columns in the list as you want them to appear in the generated file.
- Hide columns in the list that you do not want to include in the generated file. Information in hidden columns is not exported.

- Sort the rows in the list as you want them to appear in the generated file. Each list in the GUI has a column or set of columns that determine the order of its rows.
- Filter the list to show only the resources, tasks, or alerts that you want to export. Use filtering to limit the rows that are shown in a list based on value in a specific column or set of default columns.

You can export information that is shown on the following pages:

Resource list and resource details pages

You can export information about the top-level resources, internal resources, or related resources that are shown on resource list pages and resource details pages. For top-level resources, you can also export information about related tasks or alerts.

For example, you can export information about monitored switches on the Switches page. When you select the export action, the generated file includes the column values for each monitored switch in the list. In the **Alerts** tab on the Switches page, you can export information about all the alerts that are related to switches.

Alerts page

On the **Home** > **Alerts** page, you can export information about all the alerts that were detected for monitored resources. Information that you can export includes all the values that are being shown in the columns for the list of alerts.

Tasks page

On the **Home** > **Tasks** page, you can export information about all the tasks that are used to optimize resources and provision storage. Information that you can export includes all the values that are being shown in the columns for the list of tasks.

- 1. In the GUI, go to a list of resources, tasks, or alerts for which you want to export information.
- 2. Right click any row in the table and select **Export** > <**file_format>**, where <**file_format>** represents the format that you want to use for a file.

The following formats are available:

CSV (comma-separated values)

A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application. The information in a CSV file has the following format:

```
"Column_name","Column_name","Column_name","Column_name","Column_name","Column_n
ame"
"data","data","data","data","data"
```

where Column_name represents the name of a column in a table and data represents the data that is associated with a column.

PDF

You can view a PDF file with Acrobat reader. The information that you export to a PDF file is formatted into a table.

Tip: When you export a table of data that contains many columns, the rows in that table might span multiple pages in the resulting PDF. For example, if a table contains 20 columns, then the row for a specific resource might be shown on more than one page (10 columns on the first page and 10 columns on the second page).

To reduce the number of pages that a table spans, before you select the export action, hide the columns that you do not want to include in the PDF. Continue hiding columns until the table no longer spans multiple pages.

HTML

You can view an HTML file with a web browser. The information that you export to an HTML file is formatted into a table.

- 3. Optional: Depending on how your web browser is configured, you can specify the name and location of the generated file.
- 4. Click the save option to export the information to a file.

Related concepts

Viewing information about resources

You can view detailed information about the resources that are monitored by IBM Spectrum Control. Resources include storage systems, volumes, pools, servers, hypervisors, fabrics, and switches. You can also view information about internal and related resources.

Customizing lists of resources, tasks, and alerts

Customize lists to focus on the information that is important to you. You can filter information, sort rows, and show, hide, and reorder columns in the table views for resources, tasks, alerts, and other objects.

Customizing lists

Customize lists to focus on the information that is important to you. You can filter information, sort rows, and show, hide, and reorder columns in the table views for resources, tasks, alerts, and other objects.

Filtering lists

You can filter the items that are shown in a list. Use filtering when you want to limit the rows that are shown in a list based on values in a specific column or set of columns. For example, when viewing alerts, you can filter the list so that only alerts with a value of Warning in the Severity column are shown.

- 1. In the web-based GUI, go to a list of resources, jobs, or alerts that you want to filter.
- 2. Click the **Filter** icon \(\bigcup_{\text{to determine how to filter the list.} \)

The following options are available:

Filter by Default Columns

Select this option to apply the filter to the set of the default columns in a list. Each list contains a set of default columns on which you can filter.

For example, in the lists of alerts on the **Home > Alerts** page, select this option to apply filter text to the following columns at the same time: Condition, Severity, Alert Category, Resource, Internal Resource, Total Occurrences, Alert Name.

Filter by a Specific Column

Select this option to apply filter text to a specific column in a list. For example, on the **Home** > **Alerts** page, select the **Condition** column if you want to apply filter text to values in that column only.

Restriction: Some columns in lists cannot be used for filtering and are not shown under the **Filter by a Specific Column** option.

3. In **Filter** text box *Filter*..., type the text for filtering a list.

You can type letters, whole words, partial words, and numbers. When filtering on certain columns, you can select from a list of values rather than typing in filter text.

For example, when filtering on the **Status** column, you can select one of the following values as the filter text: Error, Unreachable, Warning, Normal, Unknown.

4. Press Enter to apply the filter to the list.

Only rows with values that match or partially match the filter text are shown.

For example, on the **Home** > **Alerts** page, if you select to filter on the Severity column and type critical as the filter text, only alerts that have a severity of critical and critical - acknowledged are shown in the list.

Tip: The name of a column to which a filter applies is shown next to the Filter icon .

5. To remove the filtering for a list, click **Reset**.

Sorting lists

Each list of resources, tasks, alerts, or other objects in the GUI has a column or set of columns that determines the order of its rows. You can sort the columns in a list to organize the rows according to your requirements.

For example, the list of ports on a fabric is ordered by switch name, blade slot, and port number, in that order. You can change the sort order of these columns to organize a list of resources according to your requirements.

Tip: A visual indicator **x x** is shown next to the primary column that determines the sort order of a list.

- 1. In the GUI, go to a list of resources, tasks, alerts, or other objects that you want to sort.
- 2. Complete the following tasks to sort the rows in the list:
 - To change whether the rows in a list are sorted by ascending or descending order, click the name of the primary column in the heading row. Click the name of the column again to reverse the sort order.
 - To change the column that determines the sort order of the list, click the name of that column. For example, click the Name column to sort the list of items alphabetically, or click the **Status** column to sort the list according to the status of items.

Tip: When you click a column other than the primary sort column, that column becomes the new primary sort column. However, the previous primary column maintains its sort order (ascending or descending).

For example, if a list of resources is ordered by the **Name** column in ascending order and you click the **Status** column, the rows are sorted according to the statuses of resources. At the same time, the resources for each type of status are still sorted in ascending order by the **Name** column. There is a maximum of three columns that can be sorted upon at the same time.

3. Optional: To reset the order of the rows for a column to the default setting, right-click the heading row in the list and select **Restore Default View**.

On the **Switches** page, you can use the sorting function to group virtual or logical switches with their physical switch. First, show the Virtual and IP Address columns. Click the Virtual column to sort it in ascending alphabetical order. Then, click the IP Address column.

Showing, hiding, and reordering columns in lists

Each list in the GUI has a set of columns that contain information about resources, tasks, alerts, or other objects. You can change the order of columns, show columns that are hidden, and hide columns that are shown in these lists.

- 1. In the GUI, go to a list of resources, task, alerts, or other objects that you want to customize.
- 2. Complete the following tasks to customize the columns in the list:
 - To show or hide the columns in a list, right-click any of the column headings and select or clear the check box next to a column name.
 - To change the order of columns, click a column heading and drag it to a new position in the list.

Tip: The changes that you make to column settings are saved from session to session.

3. Optional: To reset the order and display of columns to the default setting, right-click any of the column headings and select **Restore Default View**.

Managing tasks

Tasks are used to provision storage and optimize resources in your storage environment. Use the **Tasks** page to manage all the tasks that are used by IBM Spectrum Control to provision storage and optimize resources. Use resource list and resource details pages to manage the tasks for specific resources and resource types.

Discontinued support: Cloud configuration, provisioning, and optimization are no longer supported in IBM Spectrum Control. While these features might still work in this release, it's recommended that you

use another tool for your provisioning, optimization, and volume transformation needs, when possible. For a complete list of discontinued features, see "Discontinued features in IBM Spectrum Control" on page xxx.

Tasks are created when you take any of the following actions in the GUI:

- You can provision storage to assign storage to servers or hypervisors.
- You can analyze tiering to move volumes to higher or lower tiers.
- You can transform storage to move or convert volumes.
- You can balance pools to distribute the workload of volumes across pools on the same tier.

Use the following actions to manage tasks:

- View the status of tasks in a central location.
- Identify all the tasks that are related to a specific resource or a type of resource.
- Identify all the tasks that are encountering problems when the tasks are running.
- Administer the tasks. The actions that are available depend on the type of task that you select. For example, you can run or rename a tiering-analysis task, or export provisioning task details to a text file.
- Change the number of task runs that are displayed.

Tip: A *task run* is a single invocation of a task.

Viewing tasks

Use the **Tasks** page to view all the tasks that are used by IBM Spectrum Control to provision storage and optimize resources. Use resource list and resource details pages to view the tasks for specific resources and resource types.

Viewing all tasks

Use the **Tasks** page as a central location to view the tasks that IBM Spectrum Control is using to provision storage and optimize resources. You can view all the tasks for resources that are monitored by IBM Spectrum Control.

- 1. In the menu bar in the web-based GUI, go to **Home** > **Tasks**.
 - The information about tasks is organized into columns. These columns include information about the status of tasks, the date and time of the last run for tasks, and the schedule for tasks.
- 2. Optional: View the status icons at the top of the page to view a summary of task statuses.

This summary includes the number of tasks that have a Failed, Warning, Successful, or Running status.

Tip: To view descriptions of the icons and columns on the **Tasks** page, click the **Help** icon **1** in the upper-right corner of the page.

Viewing pending tasks

Some tasks might be pending because they can be run or scheduled only by an administrator. If tasks are pending, the number of pending tasks is displayed in the title bar of the GUI.

You must be assigned to the Administrator role to view the number of pending tasks.

Pending tasks are tasks that were not run and are not scheduled to run. Pending tasks might be created by a user who is not authorized to run or schedule the task. For example, users who are assigned to the Monitor or External Application role can have permission to create provisioning tasks by using a service class. However, if the service class specifies that administrator approval is required, these users can only create the provisioning task. The provisioning task is pending until an administrator runs or schedules it.

To view pending tasks, complete the following steps:

- 1. In the title bar of the GUI, look for the pending task count. If no number appears in the title bar, there are no pending tasks.
- 2. To open a **Tasks** page that shows only the pending tasks, click the number of pending tasks in the title bar.

Viewing tasks for a type of resource

Use resource list pages to view only the tasks for a specific type of top-level resource. Each resource list page has a **Tasks** tab that lists only the tasks for that type of resource. For example, you can use the list page for storage systems to view only the tasks that are associated with storage systems.

1. In the menu bar, select a type of top-level resource. For example, to view the tasks that are related to storage systems, go to **Storage** and select the type of storage system you want to monitor.

Exception: Provisioning tasks that provision network-attached storage (NAS) file shares to servers or hypervisors are not shown on the resource list pages for servers or hypervisors. You must use the **Tasks** page to view these tasks.

2. Click the **Tasks** tab to view a list of all the tasks for the monitored resources of that type.

The information about tasks is organized into columns. These columns include information about the status of tasks, the most recent date and time when tasks ran, and the schedule for tasks.

Tip: The status icon on the Tasks tab represents the most critical status of the tasks that are associated with the monitored resources. For example, if one of the tasks for storage systems failed, the following icon is shown on the tab:



Tasks

3. Optional: View the status icons on the **Tasks** page for a summary of task statuses.

This summary includes the number of tasks that have a Failed, Warning, Successful, or Running status.

Tip: To view descriptions of the icons and columns on the Tasks page, click the Help icon 🛍



Viewing tasks for a specific resource

Use resource detail pages to view only the tasks for a specific monitored resource. Each resource detail page has a Tasks link that lists only the tasks for the associated storage system, server, hypervisor, switch, or fabric.

- 1. In the menu bar, select a type of top-level resource. For example, if you want to view the tasks for a storage system, go to **Storage** and select the type of storage system you want to monitor.
- 2. Right-click a resource in the list and select **View Details**.

The resource detail page is displayed. In the General section, the number next to Tasks represents the number of tasks that are associated with the resource. For example, Tasks (3) indicates that you can view three tasks.

The icon next to the **Tasks** link represents the most critical status of the tasks that are associated with the resource. For example, if a task failed, the following icon is shown:



3. Click **Tasks** in the **General** section to view the tasks for the resource.

The information about tasks is organized into columns. These columns include information about the status of tasks, the most recent date and time when tasks ran, and the schedule for tasks.

4. Optional: View the status icons on the **Tasks** pane to view a summary of task statuses. This summary includes the number of tasks that have a Failed, Warning, Successful, or Running status.

Tip: To view descriptions of the icons and columns on the Tasks pane, click the Help icon 🙉

Viewing task details

Use task detail pages to view detailed information about the tasks that IBM Spectrum Control uses to optimize resources and provision storage. You can use this information to troubleshoot the reason why a task failed, analyze details of a task before you implement task recommendations, or schedule future task runs.

On all the task detail pages, you can view information such as the task status, the duration and completion date of task runs, and the actions that you can take to manage the task. For tasks such as tiering analysis, balance analysis, and transform plan, you can view information about the recommendations that are generated by the most recent or earlier task runs. For provisioning tasks, you can view the status of the actions that the provisioning task completes for each of the resources that are configured in the task.

- 1. From the menu bar, go to the **Tasks** page, a resource list page for a resource type, or a resource detail page for a specific resource.
 - For example, to view the tasks that are related to storage systems, go to **Storage** and select the type of storage system you want to monitor, and click the **Tasks** tab.
- 2. Locate the task that you want to view in more detail.
- 3. Right-click the task row and select View Details.
 The task details page opens in a separate window. For example, the Tiering Analysis page opens if you select a tiering-analysis task row.
- 4. Optional: To view task logs, click **Open Logs**. You can use the information in the log file to troubleshoot any errors that might occur when a task is run.
- 5. Optional: For provisioning tasks, you can export information about the provisioning task to a text file for offline viewing. To export information, click the icon at the top of the **Provisioning** page.

Viewing task logs

Task logs include detailed information about the status, actions, and progress of a task. Status information includes informational, warning, and error messages that are related to each action in a task that is taken during processing. You can use this information to troubleshoot any errors that might occur when a task is run.

The following examples show the type of information that you can view in the task logs:

- View the number of recommendations that are generated by tiering-analysis, balance-analysis, or transform-plan task runs.
- View the options that you have when you transform storage, analayze tiering, or balancing pools. For
 example, when you create a tiering-analysis task and the options that you choose are recorded in the
 task log.
- View error messages that are related to task processing.
- 1. From the menu bar, go to the **Tasks** page, a resource list page for a resource type, or a resource detail page for a specific resource.
 - For example, to view the tasks that are related to storage systems, go to **Storage** and select the type of storage system you want to monitor, and click the **Tasks** tab.
- 2. Locate the task that has a log file that you want to view.
- 3. Right-click the task row and select **Open Logs**.
 - The **Logs** page shows log entries for the most recent task run.
- 4. Optional: On the **Logs** page, to view the log for a previous task run, select a task run from the **Select a log** list.
- 5. Optional: To view only the actions in a task run that have a Warning or Error status, select an option from the **Show all** list.

You can choose to view only actions that have the following statuses:

- Only error entries
- 🖺 Only warning entries

- 6. Optional: To view an explanation of the message that is associated with an action, click the link in the **ID** column.

The information on the **Logs** page is automatically updated every 30 seconds. New entries are added to the end of a log. You can view the following information on the **Logs** page:

- The overall status of a task. The icon that is shown in the **Select a log** list represents the most critical status that was generated by an action in the task run.
- · The status for each action in a task.
- The date and time when an action was completed. The date, time, and time zone of the action is shown in the Date and Time column.
- The ID of the message that is associated with an action. You can click the value in the ID column to view more information about a message.
- The description of an action.

Tip: You can use the navigation and search functions of the web browser to locate information on the page.

Setting the number of task runs that are displayed

You can change the number of task runs that are shown in the GUI by configuring the settings on the **History Retention** page. By default, the information for the last five runs of a task is shown.

- 1. In the menu bar, go to **Settings** > **History Retention**.
- 2. Click Edit.
- 3. In the **Job logs** field, enter the number of runs that you want to retain for display.
- 4. Click **Save** to save your changes.

After you update the value for job logs, the list of task runs is automatically updated within a few minutes.

Managing tasks for provisioning

Provisioning tasks are used to create storage volumes and assign the volumes to servers or hypervisors. Provisioning tasks are also used to create network-attached storage (NAS) file shares. Use the **Tasks** page in the GUI to manage all the tasks that are used by IBM Spectrum Control to provision storage. Use resource list and resource details pages to manage the provisioning tasks for specific resources and resource types.

Discontinued support: Provisioning is no longer supported in IBM Spectrum Control. While the feature might still work in this release, it's recommended that you use another tool for your provisioning needs, when possible. For a complete list of discontinued features, see "Discontinued features in IBM Spectrum Control" on page xxx.

Use the GUI to complete the following provisioning actions:

- Create provisioning tasks
- · Run provisioning tasks immediately
- Schedule provisioning tasks to run at a specified date and time
- Export information about provisioning tasks to text files
- Delete provisioning tasks

Related concepts

"Provisioning storage overview" on page 72

IBM Spectrum Control guides you through the steps for provisioning storage volumes or network-attached storage (NAS) shares. You can provision storage volumes or NAS file shares to one or more servers, one or more hypervisors, or one cluster.

Creating provisioning tasks

You can create provisioning tasks with the GUI and assign storage volumes or NAS file shares to servers or hypervisors.

You must have Administrator privileges to create provision tasks.

Related tasks

"Provisioning storage with the IBM Spectrum Control GUI" on page 434

You can provision volumes or shares to servers, hypervisors, and clusters in IBM Spectrum Control.

"Collecting information about shares on storage systems" on page 150

To collect information about shares in file systems, you must modify the TPCD.config file. When you provision shares, you can then view such information as the name of the share, the name of the file system, and the protocols that are used to share files.

Running provisioning tasks

Use the **Tasks** page to run provisioning tasks. You can run provisioning tasks immediately or you can schedule tasks to run at a date and time that you specify. When you run a provisioning task, the configuration changes that are contained in the task are implemented. For example, storage volumes are created and assigned to servers.

To run provisioning tasks, you must be the owner of the task or have Administrator privileges.

Provisioning tasks can run only once. You can run the task when you complete the steps for storage provisioning or you can save the provisioning task and run the task later.

The **Run Now** and **Schedule** actions are available for a provisioning task when the task is not running, and was not previously run.

- 1. From the menu bar, go to the **Tasks** page, a resource list page for a resource type, or a resource detail page for a specific resource.
 - For example, to view the provisioning tasks that are related to servers, go to **Servers** > **Servers**, and click the **Tasks** tab.
- 2. Locate the provisioning task that you want to run.
- 3. Depending on when you want the provisioning task to run, choose one of the following options:
 - To run the provisioning task immediately, right-click the task row and select **Run Now**. The task status changes to Running and the provisioning task runs immediately.
 - To schedule the provisioning task to run at a specific date and time, complete the following steps:
 - a. Right-click the task row and select **Schedule**.
 - b. From the **Schedule** page, select a future date and time from the **Provision** lists, and click **Save**.

You can also run or schedule the provisioning task from the **Provisioning** detail page. On the **Provisioning** detail page, click **Execute** to run the task immediately or click **Schedule** to schedule the task to run at a specific date and time.

The provisioning task runs immediately or on the scheduled date and time. To check the progress of the task, you can take the following actions:

- View the task status on the **Tasks** page.
- View detailed information about the progress of the task on the **Provisioning** detail page. You can view the task duration and the status of the actions that the provisioning task completes for each of the resources that are configured in the task.
- View the task logs from the **Tasks** page or from the **Provisioning** detail page. You can view detailed informational, warning, and error messages that are related to the task. You can use this information to troubleshoot any errors that might occur when a task is run.

Exporting provisioning task information

Use the **Provisioning** detail page to export details of a provisioning task to a text file. The task details that are shown on the **Provisioning** page are saved to the generated text file.

If you do not have authority to provision storage, you can export the provisioning task details to a file and email the file to a user with Administrator privileges. The administrator can run or schedule the provisioning task if they approve the task details.

- 1. From the menu bar, go to the **Tasks** page, a resource list page for a resource type, or a resource detail page for a specific resource.
 - For example, to view the provisioning tasks that are related to servers, go to **Servers** > **Servers**, and click the **Tasks** tab.
- 2. Locate the provisioning task that you want to export.
- 3. Right-click the task row and select View Details.
- 4. From the **Provisioning** page, to start the export, click the 🖃 icon in the upper-right corner of the page.
- 5. Save the generated text file.

Depending on how your web browser is configured, you can specify the file name and the location where the file is saved.

The provisioning task information is exported to a text file. The generated file includes all the provisioning task details that are listed on the **Provisioning** page. Columns that are hidden on the **Provisioning** page are included in the generated file. The information in the file is organized according to a default order and is not affected by the column order that is defined on the **Provisioning** page.

Managing tasks for tiering storage, balancing pools, and transforming storage

To optimize the resources in your storage environment, use tiering storage, balancing pools, and transforming-storage tasks. Use the **Tasks** page to manage all the tasks that are used to optimize storage. Use resource list and resource detail pages to manage the optimization tasks for resource types and for specific resources.

Discontinued support: Storage optimization is no longer supported in IBM Spectrum Control. While the feature might still work in this release, it's recommended that you use another tool for your balancing and tiering needs, when possible. The storage transformation feature is still available. For a complete list of discontinued features, see "Discontinued features in IBM Spectrum Control" on page xxx.

To schedule, run, edit, or delete tiering storage, balancing pools, and transforming-storage tasks, you must have Administrator privileges.

To complete any of the optimization tasks, such as balancing pools, you must complete the following steps:

- 1. Create the analysis task with the GUI. For example, to balance pools create a balance-analysis task.
- 2. Run the analysis task to generate the optimization recommendations. For example, run the balance-analysis task to generate the recommendations to move volumes.
- 3. Edit analyze tiering and balancing pool tasks by modifying the thresholds.
- 4. Implement the recommendations to optimize storage. For example, implement the recommendations to balance pools by running the analysis-execution task.

Use the **Tasks**, resource list, and resource detail pages to take the following actions for optimization tasks:

- Run the optimization analysis immediately to generate optimization recommendations.
- Schedule the optimization analysis to run at a specified date and time.
- Implement optimization recommendations immediately.
- Schedule the implementation of optimization recommendations for a specified date and time.
- Pause, cancel, and resume the implementation of optimization recommendations.
- Delete optimization tasks.

· View task logs.

Related tasks

Modifying the criteria for analyzing tiering

When a task is created after you run the tiering analysis, you can modify the criteria for tiering the volumes.

Modifying the criteria for balancing pools

When a task is created after you run the analysis for balancing pools, you can modify the criteria that determine which pools are balanced.

Creating analysis tasks

To create tasks for optimizing resources in your storage environment, use IBM Spectrum Control GUI.

You must have Administrator privileges to create tasks for optimizing resources in your storage environment.

Storage optimization tasks are created when you take any of the following actions in the GUI:

Tier storage

You can analyze tiering to move volumes to higher or lower tiers; a tiering-analysis task is created.

Balance pools

You can balance pools to distribute the workload of volumes across pools on the same tier; a balance-analysis task is created.

Transform storage

You can transform storage to move or convert volumes; a transform-plan task is created.

Related tasks

"Tiering volumes by I/O density and I/O rate" on page 443

Set thresholds for tiering volumes by I/O rate or I/O density. To gain better optimization results, you can also customize the analysis of the workload activity of the pools.

"Transforming and migrating volumes" on page 448

You can transform fully allocated volumes to compressed or thin-provisioned volumes, or transform compressed or thin-provisioned volumes to fully allocated volumes. You can move volumes to other pools or to pools that are enabled for Easy Tier.

Running analysis tasks

To run tiering analysis, balance analysis, and transform-plan tasks, use the **Tasks** page. When you run the analysis for any of these optimization tasks, recommendations are generated to optimize the storage resources that are configured in the task. For example, when you run a tiering-analysis task, recommendations are generated to tier volumes based on the criteria that you set in tiering policies.

To run tiering analysis, balance analysis, and transform-plan tasks, you must have Administrator privileges.

You can run tasks immediately or schedule tasks to run at a future date and time.

Tip: The Run Now action is available when the task is enabled and the task is not running.

- 1. From the menu bar, go to the **Tasks** page, a resource list page for a resource type, or a resource detail page for a specific resource.
 - For example, to view the tiering-analysis tasks that are related to servers, go to **Servers** > **Servers**, and click the **Tasks** tab.
- 2. Locate the tiering analysis, balance analysis, or transform-plan task that you want to run.
- 3. Right-click the task row and select Run Now.
 - The task status changes to Running.
- 4. Optional: To schedule future task runs for tiering analysis or balance-analysis tasks, right-click the task row and select **Schedule Analysis**. On the **Schedule Analysis** page, select the time that the task is run and how often the task is run.

Restriction: You cannot schedule future task runs for transform-plan tasks.

5. Optional: To enable or disable a tiering analysis or balance-analysis task, select **Enabled** or **Disabled** on the **Schedule Analysis** page.

When the task is enabled, the task runs according to the defined schedule. If the task is disabled, you can modify the task schedule but the task cannot run at the scheduled time.

When the task run completes, you can view the recommendations that are generated on the task detail page. For example, you can view the recommendations that are generated for a tiering-analysis task on the **Tiering Analysis** page. To check the progress of the task run, you can take the following actions:

- View the task status on the **Tasks** page.
- View information about the progress of the task on the task detail page. You can view the duration of the task run, and the date and time that the task run completed.
- View the task logs from the Tasks page or from the task detail page. You can view detailed
 informational, warning, and error messages that are related to the task. You can use this information to
 troubleshoot any errors that might occur when a task is run.

Related tasks

"Modifying the criteria for analyzing tiering" on page 444

When a task is created after you run the tiering analysis, you can modify the criteria for tiering the volumes.

Implementing recommendations to optimize storage

Use the **Tasks** page and task detail pages to implement optimization recommendations to tier storage, balance pools, and transform storage. You can implement recommendations immediately or you can schedule the implementation of the recommendations for a future date and time.

Related tasks

"Modifying the criteria for analyzing tiering" on page 444

When a task is created after you run the tiering analysis, you can modify the criteria for tiering the volumes.

Implementing recommendations immediately

To implement optimization recommendations to tier storage, balance pools, and transform storage, use the task detail pages. For example, use the **Balance Analysis** page to implement the recommendations to balance pools.

To implement storage optimization recommendations, you must have Administrator privileges.

IBM Spectrum Control uses analysis-execution tasks to implement optimization recommendations. When you implement recommendations for a tiering analysis, balance analysis, or transform-plan task, the following process takes place:

- An analysis-execution task is automatically created so that the recommendations can be implemented.
- The analysis-execution task runs immediately and the optimization recommendations are implemented. For example, volumes are moved to higher or lower tiers when storage-tiering recommendations are implemented.

Tip: In certain instances, running several optimization recommendations at the same time can cause one or more of the individual recommendation executions to encounter a Pool Exhaustion Exception. In these instances, the executions error out, rather than waiting for resources in the thread pool to become free. You might need to manually finish the execution of some recommendations, in particular, the old VDisk copy might not be removed, and you must remove it manually. To reduce the likelihood of encountering the issue again, contact IBM Support for assistance in decreasing the total number of concurrent executions allowed.

You can view analysis-execution tasks on the **Tasks** page. The Related Task column shows the related tiering analysis, balance analysis, or transform-plan task.

The **Execute** action is available under the following conditions:

• Recommendations were generated by the latest run of the tiering analysis, balance analysis, or transform-plan task.

- An analysis-execution task was not already run for the latest set of recommendations.
- An analysis-execution task is not running or paused.
- 1. From the menu bar, go to the **Tasks** page, a resource list page for a resource type, or a resource detail page for a specific resource.
 - For example, to view the analysis-execution tasks that are related to servers, go to **Servers** > **Servers**, and click the **Tasks** tab.
- 2. Locate the tiering analysis, balance analysis, or transform-plan task that you want to implement.
- 3. Right-click the task row and select **View Details**.

 For example, the **Balance Analysis** page opens if you select a balance-analysis task.
- 4. Click Execute.

An analysis-execution task is created and immediately implemented.

The analysis-execution task runs immediately. To check the progress of the task, you can take the following actions:

- View the task status on the **Tasks** page.
- View the status of individual task recommendations on the Analysis Execution page.
- View the task logs from the **Tasks** page or from the **Analysis Execution** page. You can view detailed informational, warning, and error messages that are related to the task. You can use this information to troubleshoot any errors that might occur when a task is run.

Scheduling the implementation of recommendations

To schedule the implementation of optimization recommendations to tier storage, balance pools, and transform storage, use the **Tasks** page.

To schedule the implementation of storage optimization recommendations, you must have Administrator privileges.

The **Schedule Execution** action is available under the following conditions:

- Recommendations were generated by the latest run of the tiering analysis, balance analysis, or transform-plan task.
- An analysis-execution task is not running or paused.
- An analysis-execution task was not already run for the latest set of analysis recommendations. For each run of the tiering analysis, balance analysis, or transform-plan task, you can complete only one run of the analysis-execution task.

When you define the schedule, an analysis-execution task is created for the implementation of the optimization recommendations.

Tip: The analysis-execution task implements the recommendations that were current when the analysis-execution task was defined. For example, you define a tiering-analysis task that runs daily. You review the recommendations every day and, on a Tuesday, you schedule an analysis-execution task to run on the following Friday. When the analysis-execution task runs on Friday, the recommendations that were current on Tuesday are implemented, even if more recent recommendations are available.

- 1. From the menu bar, go to the **Tasks** page, a resource list page for a resource type, or a resource detail page for a specific resource.
 - For example, to view the tiering analysis tasks that are related to servers, go to **Servers** > **Servers**, and click the **Tasks** tab.
- 2. Locate the tiering analysis, balance analysis, or transform-plan task that you want to schedule.
- 3. Right-click the task row and select **Schedule Execution**.
- 4. On the **Schedule Execution** page, select the date and time that the recommendations are implemented and click **Save**.

An analysis-execution task is created that you can monitor on the **Tasks** page. The name of the analysis-execution task is based on the name of the related tiering analysis, balance analysis, or transform-plan

task. The Related Task column shows the related tiering analysis, balance analysis, or transform-plan task.

When the analysis-execution task runs on the scheduled date, the optimization recommendations are implemented. For example, volumes are moved to higher or lower tiers when storage-tiering recommendations are implemented.

Tip: In certain instances, running several optimization recommendations at the same time can cause one or more of the individual recommendation executions to encounter a Pool Exhaustion Exception. In these instances, the executions error out, rather than waiting for resources in the thread pool to become free. You might need to manually finish the execution of some recommendations, in particular, the old VDisk copy might not be removed, and you must remove it manually. To reduce the likelihood of encountering the issue again, contact IBM Support for assistance in decreasing the total number of concurrent executions allowed.

Pausing, resuming, and canceling the implementation of recommendations

IBM Spectrum Control uses analysis-execution tasks to implement optimization recommendations to tier storage, balance pools, and transform storage. To pause, resume, and cancel analysis-execution tasks, use the **Tasks** page.

To pause, resume, or cancel analysis-execution tasks, you must have Administrator privileges.

- 1. From the menu bar, go to the **Tasks** page, a resource list page for a resource type, or a resource detail page for a specific resource.
 - For example, to view the analysis-execution tasks that are related to servers, click **Servers** > **Servers**, and click the **Tasks** tab.
- 2. Locate the analysis-execution task that you want to pause, resume, or cancel.
- 3. Right-click the task row and select Pause, Resume, or Cancel.

The Pause action is only available when the analysis-execution task is running.

The **Resume** action is only available when the analysis-execution task is paused.

The **Cancel** action is available when the task is running or paused.

When you pause, resume, or cancel analysis-execution tasks, the following processing occurs:

Pause

The status of the analysis-execution task changes to Paused and the recommendations that are being implemented are suspended.

Resume

The status of the analysis-execution task changes to Running and the recommendations are resumed.

Cancel

The status of the analysis-execution task changes to Canceled and the recommendations that are being implemented are stopped. Any recommendations that are already implemented for this analysis-execution task are not rolled back. For example, a recommendation is generated to move a volume from pool_a to pool_b. If this recommendation is already implemented before the task is canceled, the volume remains in pool_b. Use the **Analysis Execution** page to determine which recommendations were implemented for the task before the task was canceled.

You cannot restart an analysis-execution task that is canceled. If you want to implement the same set of recommendations again, you must complete the following steps:

- 1. Delete the analysis-execution task that you canceled.
- 2. Use the **Schedule Execution** action for the tiering analysis, balance analysis, or transform-plan task to create an analysis-execution task.

952

Renaming tasks

You can change the name of a task by using the task detail page. For example, use the **Provisioning** page to change the name of a provisioning task. A task name is automatically generated when the task is created. If the original name is ambiguous or is not clear enough, you can change the task name to a unique task name that you provide.

You must have the following privileges to rename the different task types:

Provisioning

To rename provisioning tasks, you must be the owner of the task or have Administrator privileges.

Other task types

To rename all other tasks, such as tiering analysis, balance analysis, transform plan, and analysis-execution tasks, you must have Administrator privileges.

1. From the menu bar, go to the **Tasks** page, a resource list page for a resource type, or a resource detail page for a specific resource.

For example, to view the tasks that are related to servers, go to **Servers** > **Servers**, and click the **Tasks** tab.

- 2. Locate the task that you want to rename.
- 3. Right-click the task row and select View Details.

The task detail page opens in a separate window. For example, the **Provisioning** page opens if you select a provisioning task.

4. Type the new name in the task name field.

The naming conventions for task names are:

- The maximum length is 64 characters.
- Uppercase and lowercase letters are allowed.
- Spaces and numbers are allowed.
- The following special characters are allowed:

5. Click Rename.

The task is renamed. You can view the renamed task on the **Tasks** page.

Deleting tasks

Use the **Tasks** page to delete tasks, such as provisioning, tiering analysis, balance analysis, transform plan, and analysis-execution tasks.

You must have the following privileges to delete the different task types:

Provisioning

To delete provisioning tasks, you must be the owner of the task or have Administrator privileges.

Other task types

To delete all other tasks, such as tiering analysis, balance analysis, transform plan, and analysis-execution tasks, you must have Administrator privileges.

The following conditions apply for the different task types:

Analysis-execution

You can delete analysis-execution tasks that were canceled, already run, or scheduled to run in the future. You cannot delete tasks that are running or paused.

Other task types

You can delete all other tasks if the tasks were already run or if the tasks are scheduled to run in the future. For example, you can delete provisioning, tiering analysis, balance analysis, and transformplan tasks that were already run. You cannot delete tasks that are running.

- 1. From the menu bar, go to the **Tasks** page, a resource list page for a resource type, or a resource detail page for a specific resource.
 - For example, to view tasks that are related to servers, go to Servers > Servers, and click the Tasks tab.
- 2. Locate the task that you want to delete.
- 3. Right-click the task row and select **Delete**. Alternatively, you can delete the task by clicking **Delete** on the task detail page.
 - The **Delete** page opens in a separate window.
- 4. If you are deleting a provisioning task that was not yet run, enter the reason that you are deleting the task.
- 5. Click Delete.

To exit the window without deleting the task, click **Cancel**.

Planning copy data resources

To plan for your block storage requirements, you need to be able to see how you currently use block storage in your storage environment.

You can use the Copy Data view the following aspects of your block storage environment:

- You can see how much of your current allocated capacity is taken up by both the primary and backup data.
- You can see the ratio of allocated capacity that is taken up by primary data and plan for backup data according to your security policies.
- You can analyze storage device relationships to identify how redundancy is affecting available capacity.
- You can view how each type of storage you use is deployed in your environment.
- You can identify the volumes that are not backed up.

Monitoring copy data resources

You can monitor the use and type of block storage resources in Copy Data.

You can get a complete breakdown of how each type of storage is used across the block storage environment.

Remote Relationships

Detailed information on remote relationships includes the type of relationship that exists between the source and target storage systems.



Synchronous relationship

Host writes are delayed until the source receives confirmation that the data is written to the target volume.



Asynchronous relationship

Host writes can continue while data is being written to the target volume.

Where the type of volume is known, one of the following icons appears next to the volume in question.



Standard volume

A volume that is not thin-provisioned, compressed, encrypted, or encrypt-able.



Thin provisioned volume

The volume is only allocated the required amount of server space as required.



Compressed VDisk volume

Data is compressed as it is written to the VDisk.

Consistency Group

Details the status and type of the consistency group, and the number of relationships that exist in the group.

FlashCopy

Details the information on the last FlashCopy operation between each source and target storage system.

VDisk Mirrors

Details the information on the volumes, pools, tiers, and synchronization status of the VDisk mirrors.

- Identify target and source mirror disks at a glance.
- Easily see the disks that are compressed, thin, or normal.

HyperSwap

Details the information on the status of the master and auxiliary volumes, pools, and tiers of the HyperSwap relationships.

- Quickly identify the master and auxiliary volumes and pools.
- Drill down to see the amount of allocated disk for each volume.

Unprotected Volumes

Details the information on the volumes that are not backed up.

IBM Spectrum Control REST API

You can use the Representational State Transfer (REST) API for IBM Spectrum Control to access information about resources and to generate custom capacity, configuration, and performance reports.

To get the information that you need about your resources, you can use a REST command line utility or you can use a web browser.

Retrieving data about resources by using a REST API command line utility

You can connect to the REST API for IBM Spectrum Control and retrieve data by using a REST command line utility.

You can use any REST command line utility, for example, GNU Wget. To access information about the GNU Wget utility, go to https://www.gnu.org/software/wget/.

The IBM Spectrum Control REST API is hosted here: https://<hostname>:9569/srm/REST/api/v1/.

You can enter a command to authenticate with the REST API and store a security token in a file to use in your later queries, for example, cookies.txt.

You can use the utility, by entering the following command:

```
wget --post-data "j_username=<user name>&j_password=<password>"
--no-check-certificate --keep-session-cookies --save-cookies cookies.txt
https://<hostname>:9569/srm/j_security_check
```

You can use the security token in all subsequent commands that are issued against the REST API, by entering the following command:

```
wget --no-check-certificate --load-cookies cookies.txt
https://<hostname>:9569/srm/REST/api/v1/
```

Example (partial listing)

```
},

{
    "Description": "Provides a list of ports belonging to a parent storage system.
Format: \/StorageSystems\/<id>\/Ports For a list of ports associated with a
port owner, specify the owner type and id in the URL.
```

To see a list of all storage systems, you can append the URL:

```
https://<hostname>:9569/srm/REST/api/v1/StorageSystems
```

To see more information about a resource type, you can further append the URL (where *<Name>* is the resource type, such as StorageSystems, Switches, Servers, etc.)

```
https://<hostname>:9569/srm/REST/api/v1/<Name>
```

To see more information about a specific resource, add the ID to the URL (where $\langle id \rangle$ is available from https://<hostname>:9569/srm/REST/api/v1/ $\langle Name \rangle$):

```
https://<hostname>:9569/srm/REST/api/v1/<Name>/<id>
```

To query for volumes that belong to a specific storage system, enter:

```
https://<hostname>:9569/srm/REST/api/v1/StorageSystems/57909/Volumes
```

Example (partial listing)

```
Г
    £
          "Acknowledged": "No"
         "Allocated Space": "1.00",
"Capacity": "1.00",
"Controller": "Node 0",
         "Copy Relationship": "",
"Easy Tier": "Tiered Pools\/No",
"Encryption": "",
         "Enterprise HDD Capacity": "1.00",
"Format": "FB",
"Hosts": "18",
"LSS or LCU": "00",
         "Last Data Collection": "Apr 10, 2016, 15:03:57",
"Name": "CET_RA_vol",
"Nearline HDD Capacity": "",
          "Physical Allocation": "Fully Allocated",
"Pool": "General_Use_1",
          "RAID Level": "RAID 5",
"SSD Capacity": "",
         "Service Class": ""
"Shortfall": "",
          "Status": "Normal",
"Storage System": "DS8000-2107-75BLG91-IBM",
          "Storage Virtualizer": "None",
"Thin Provisioned": "No",
          "Ticket": ""
          "Tier Distribution": "0"
         "Unallocated Space": "0.00",
"Unused Space": ""
          "Unused Space":
          "Used Allocated Space": "",
          "Used Space": "1.00"
         "Virtualizer Disk": "None",
"Volume ID": "00000",
          "Volume Number": "0",
"Volume Unique ID": "75blg91\/0000",
          "id": "79852"
     },
```

The information that is provided in the Format: field in the example shows what combinations are possible. For example,

https://<hostname>:9569/srm/REST/api/v1/StorageSystems/<id>/Ports

Tips:

- If the Format: field is omitted, you can add the ID of one resource to the URL.
- The name elements in the URL are case-sensitive.
- The names of resource types are plural.
- If the URL is not used correctly, you might receive the following error message:

```
{"result":{"type":"E","msgId":"BPCUI0099E","time":"Apr 4, 2016 16:25:07","text": "The storage resource is not available."}}
```

• If the URL returns a blank page, it means that no data is available for the resource. For example, if you query an application for file shares and there are no file shares in the application, the resulting page is blank.

Table 222. IBM Spectrum Control REST API URL examples		
IBM Spectrum Control REST API	Description	
https:// <i><hostname></hostname></i> :9569/srm/ REST/api/v1/StorageSystems	A list of all monitored storage systems.	
https:// <hostname>:9569/srm/ REST/api/v1/StorageSystems/2002</hostname>	Specific information for one storage system (ID is available from this URL: https:// <hostname>:9569/srm/REST/api/v1/ StorageSystems).</hostname>	
https:// <hostname>:9569/srm/ REST/api/v1/StorageSystems/2002/ Volumes</hostname>	A list of all volumes for one storage system.	
https:// <hostname>:9569/srm/ REST/api/v1/StorageSystems/2002/ Volumes/2465/Performance</hostname>	The available performance metrics for a volume.	
https:// <hostname>:9569/srm/ REST/api/v1/StorageSystems/2002/ RemoteReplication</hostname>	The remote mirror replication relationships for storage systems.	
https:// <hostname>:9569/srm/ REST/api/v1/StorageSystems/2002/ Volumes/2465/Performance/824? granularity=sample&startTime=145581840 0000&endTime=1455904800000</hostname>	Sample performance granularity information and a start and end time for the performance data.	

Note: IBM Storage is changing the capacity terminology that is used in IBM Storage products to make it more consistent. The new capacity terminology is used in IBM Spectrum Control Version 5.3.6 and later. The new terminology is used in the GUI, but the names used in the REST API have not changed. The output from the REST API continues to use Allocated Space instead of Used Capacity, Physical Allocation instead of Used Capacity %, and so on.

Retrieving data by using REST API with a web browser

A convenient way to use the IBM Spectrum Control REST API is within the interface itself. You can quickly access resource information for your reporting capabilities by using a web browser.

To use the REST API with a web browser, use these steps:

1. Log on to the IBM Spectrum Control GUI. To start IBM Spectrum Control, see <u>"Starting IBM Spectrum Control"</u> Control" on page 44.

- 2. Open a new tab in your web browser.
- 3. Enter https://<hostname>:9569/srm/REST/api/v1/.

For example, to see a list of fabrics, enter https://<hostname>:9569/srm/REST/api/v1/Fabrics in a new tab of your web browser (partial listing).

```
"Acknowledged": "No",
    "Active Zone Set": "",
    "Connected Switch Ports": "0",
    "Custom Tag 1": "",
    "Custom Tag 2": "",
    "Custom Tag 3": "",
    "Data Source Count": "1",
    "Fabric Type": "Cisco",
    "Last Successful Probe": "N\/A",
    "Links": "0",
    "Location": "",
    "NPV Switches": "0",
    "Name": "VSAN0002",
    "Parent Fabric": "san_director",
    "Principal Switch of Fabric": "tpc-70swt-csc",
    "Probe Schedule": "N\/A",
    "Probe Status": "N\/A",
    "Status": "Error",
    "Switch Ports": "0",
    "Switch Ports": "0",
    "Switches": "1",
    "Virtual": "Yes",
    "WWN": "2002000DECAC5081",
    "id": "2663"
```

To see the information for a specific fabric in your environment, append the URL with the ID:

```
https://<hostname>:9569/srm/REST/api/v1/Fabrics/2017
```

Example (partial listing)

```
"Acknowledged": "No",
   "Active Zone Set": "",
   "Connected Switch Ports": "18",
   "Custom Tag 1": "",
   "Custom Tag 2": "",
   "Custom Tag 3": "",
   "Data Source Count": "2",
   "Fabric Type": "Cisco",
   "Last Successful Probe": "N\/A",
   "Links": "4",
   "Location": "",
   "NPV Switches": "0",
   "Name": "san_director",
   "Parent Fabric": "",
   "Principal Switch of Fabric": "",
   "Probe Schedule": "N\/A",
   "Probe Status": "N\/A",
   "Status": "Error",
   "Switch Ports": "104",
   "Switches": "2",
   "Virtual": "No",
   "WWN": "",
   "id": "2017"
```

To see the switches that belong to the fabric "id": "2017", enter:

```
https://<hostname>:9569/srm/REST/api/v1/Fabrics/2017/Switches
```

Example (partial listing)

```
"Custom Tag 2": "",
    "Custom Tag 3": "",
    "Data Source Count": "1",
    "Domain ID": "10",
    "Fabric": "unstable BRCD_1-11",
    "Firmware": "v6.4.3g",
    "IP Address": "9.11.91.241",
    "Last Successful Monitor": "None",
    "Last Successful Probe": "None",
    "Last Successful Probe": "None",
    "Location": "",
    "Mode": "Native",
    "Model": "Brocade 4100",
    "Name": "mdm-y76-swt",
    "Parent Switch": "",
    "Performance Monitor Interval (min)": "",
    "Performance Monitor Status": "Disabled",
    "Ports": "",
    "Principal Switch of Fabric": "cjswitch4",
    "Probe Status": "Never Probed",
    "Serial Number": "1070029",
    "Status": "Normal",
    "Vendor": "IBM",
    "Virtual": "No",
    "WWN": "1000000051E347790",
    "id": "60989"
```

Note: IBM Spectrum Control can connect to Brocade switches and fabrics either directly or by using Brocade Network Advisor. For Brocade switches and fabrics that IBM Spectrum Control connects to directly, the Data Source Count value is 0.

Table 223. IBM Spectrum Control REST API URL examples		
IBM Spectrum Control REST API	Description	
https:// <i><hostname></hostname></i> :9569/srm/ REST/api/v1/StorageSystems	A list of all monitored storage systems.	
https:// <hostname>:9569/srm/ REST/api/v1/StorageSystems/2002</hostname>	Specific information for one storage system (ID is available from this URL: https:// <hostname>:9569/srm/REST/api/v1/ StorageSystems).</hostname>	
https:// <hostname>:9569/srm/ REST/api/v1/StorageSystems/2002/ Volumes</hostname>	A list of all volumes for one storage system.	
https:// <hostname>:9569/srm/ REST/api/v1/StorageSystems/2002/ Volumes/2465/Performance</hostname>	The available performance metrics for a volume.	
https:// <hostname>:9569/srm/ REST/api/v1/StorageSystems/2002/ RemoteReplication</hostname>	The remote mirror replication relationships for storage systems.	
https:// <hostname>:9569/srm/ REST/api/v1/StorageSystems/2002/ Volumes/2465/Performance/824? granularity=sample&startTime=145581840 0000&endTime=1455904800000</hostname>	Sample performance granularity information and a start and end time for the performance data.	

Note: IBM Storage is changing the capacity terminology that is used in IBM Storage products to make it more consistent. The new capacity terminology is used in IBM Spectrum Control Version 5.3.6 and later. The new terminology is used in the GUI, but the names used in the REST API have not changed. The output from the REST API continues to use Allocated Space instead of Used Capacity, Physical Allocation instead of Used Capacity %, and so on.

Chapter 7. Reporting

Keep your colleagues and management up-to-date by sending inventory, capacity, performance, and storage consumption reports by email or by exporting information about your storage environment.

Reports overview

Share the information that is collected about your storage environment by sending reports by email or by exporting the information to create your own reports.

In IBM Spectrum Control, you can share information with your colleagues and management:

- By creating predefined inventory and capacity reports that you can save to your file system and send by email
- By creating custom reports that you can save to your file system or send by email
- By creating chargeback and consumer reports that you can send by email
- By using the Representational State Transfer (REST) API to get information about storage resources and to generate configuration and performance reports
- By exporting information about the configuration, capacity, and performance of storage resources to HTML, PDF, or CSV files

More about exporting information to files:


You can use the export feature as follows:

- To export capacity and performance information about your storage resources from IBM Spectrum Control into spreadsheets and the reports that you create for your organization
- To share information with your colleagues about the storage resources in your environment
- To make your colleagues aware of issues and to help them investigate and resolve capacity and performance issues

You can use an external reporting tool, Cognos Analytics reporting interface, to view predefined reports and create custom reports about IBM Spectrum Control.

Predefined, custom, and storage consumption reports


Reports



Leverage and share the insights that IBM Spectrum Control gains from the metadata that it collects and analyzes to generate information about your storage assets and to monitor capacity usage and plan capacity purchases.

- Create predefined capacity reports
- Create predefined inventory reports

Custom reports



Create custom reports

Chargeback and consumer reports

Create, schedule, and send chargeback and consumer reports by email to make your organization aware of the cost and the amount of capacity that is used by storage consumers such as applications.

- More information about chargeback and consumer reports
- Create chargeback reports
- Create consumer reports
- 1. Creating predefined capacity reports
- 2. "Creating predefined inventory reports" on page 482
- 3. "Creating custom reports" on page 484
- 4. Storage consumption reports: Overview
- 5. "Creating chargeback reports" on page 494
- 6. "Creating consumer reports" on page 496

In the IBM Spectrum Control you can create the following types of reports, which you can configure, schedule, and send by email:

Predefined reports

You can use predefined reports:

- To notify colleagues about the capacity issues that affect them, and to monitor the capacity usage of resources, such as storage systems, pools, and servers, to ensure that you have sufficient capacity to meet current usage and to help plan capacity purchases
- To generate inventory reports about your storage assets, such as storage systems, hypervisors, servers, IBM Spectrum Virtualize nodes, and IBM Spectrum Scale nodes

Custom reports

You use custom reports to create a variety of reports such as reports that provide information about your storage assets, and capacity and performance reports that provide information about the storage usage and performance of your resources.

Tip: You can open the table view for any resource, such as storage systems, volumes, pools, general groups, applications, and departments and click **Create Report** to create a custom report that you can save or send in a few simple steps. To create performance reports, open the **Performance** page for the resource and click **Create Report**.



Storage consumption reports

Chargeback reports make the owners of resources, such as applications, departments, hypervisors, or physical servers, aware of the amount and cost of the block and file storage that is consumed.

Consumer reports make the owner of a specific resource, such as an application, department, hypervisor, or physical server, aware of the amount and cost of the block and file storage that is consumed.

How to create reports

Watch a short video about how to create and customize a capacity report, attach a file to the report, and schedule when to run it in IBM Storage Insights Pro.

Creating predefined capacity reports

To share information about capacity availability, usage, and growth, create capacity reports. You can configure and refine the information that is included in the report so that your colleagues get the information that they need to monitor the resources that they manage and make capacity planning decisions.

Before you start

If you want to send reports by email, you must set up the mail server. If you already set up the mail server to send alerts or reports, or if you want to only save reports to your file system, you don't have to set up the mail server.

If you want the reports that you will create to be saved to your file system to be stored in a folder other than the default reports folder, you must create the folder before you create the reports.

Saving reports to file systems on remote servers: The machine on which the Web server component was installed and runs must be able to access the remote or shared file system that is specified for saving reports. To make file systems on a server available to remote clients, they must be mounted as a CIFS export on Windows operating systems or as an NFS export on UNIX or Linux operating systems. On Windows operating systems, the file system must be mounted as the local system user or from a service that is running under the local system account so that it is visible to all logon sessions. One way of mapping a drive as the local system user is to map the drive from a psexec command window. The psexec command window is available as part of the Sysinternals command-line tools for Windows operating systems.

Step 1: Pick your predefined report

Click **Reports**, and then click **Create Report** and pick the capacity report that you want to create.


Step 2: Name the report

Provide a unique name for your report. You can use alphabetical and numerical characters, hyphens, dashes, and blank spaces.

Step 3: Specify the scope of the report

Choose one of these options:

- Option 1: Generate capacity information about all of the resources of a specific type.
- Option 2: Generate capacity information about a selection of resources of a specific type.
- Option 3: Generate capacity information about the resources of a specific type in a group.



Option 1: All of the resources


Choose this option to generate a report about all of the resources, such as all storage systems, pools, volumes, managed disks, or NAS file systems.

Option 2: One or more of the resources

Choose this option to generate capacity information about a selection of resources such as a selection of storage systems, pools, volumes, managed disks, or NAS file systems.

When you create capacity reports about managed disks, pools, servers, or volumes, you can group the types of resources, such as managed disks by storage systems, pools by storage systems, servers by applications, and volumes by servers.

Try it out: Click **Reports**, click **Create Report**, and choose the **Pools** capacity report. Choose capacity information for one or more resources, and select pools by storage systems or pools by tiers. To refine the list of the resources that are shown, you can enter a matching pattern for the names of the storage systems, or you can select the storage systems that contain the pools that you want to include in the report.



Option 3: Resources assigned to groups

You can also generate capacity information about the storage systems that belong to a group that you created, such as a general group that consists of the storage systems that are assigned to your production platform, or the storage resources that are assigned to an application or department.

Tip: The report you pick, determines the level of information that is shown about applications in reports. For example, if you want to see the capacity information about the set of volumes that you added to an application, create a **Volumes** capacity report and select the application. Or, if you want to see the capacity information about the filesets that you added to an application, create a **Filesets** capacity report and select the application.


Step 4: Select columns

To ensure that the report's recipients get the information that they need, you can add more information, such as the location of the resources or custom tags. You can also remove the information that your colleagues don't need.

Step 5: Add filters

To refine the information in the report, you can add up to four filters.

Sample filters for block storage systems capacity reports



For example, if you want to notify resource administrators or owners about capacity shortfalls, you can set threshold values for capacity information. Or if you want to exclude information that is not of interest to the report's recipients, you can add a filter.

Don't forget to check the table when you apply filters to make sure that the recipients of the report will get the information that they need.

Tip: You can add advanced filters to refine information about capacity and information about the assets in your storage environment. If advanced filters aren't available for a resource, you can use the filter feature that is available in the table. For example, in reports about volumes, you can filter by volume, pool, or storage system name, or by Easy Tier attributes.

Step 6: Schedule and deliver the report

To notify your colleagues of capacity or configuration anomalies, create one-off capacity and inventory reports or create a schedule to provide regular updates about the capacity or state of your storage assets.

Tip: Instead of sending the report by email, you can save the report to the default reports folder or to a folder on your file system. Alternatively, when you send the report by email, you can also attach the report as a CSV, PDF, or HTML file.

Review, edit, and delete reports

To review the report that you created, click **Reports**, expand the relevant section, such as **Capacity Report** for predefined capacity reports and **Inventory Report** for predefined inventory reports, and then select the report.

You can edit the report, such as change the name and scope of the report or schedule of the report. To delete the report, click **Actions** > **Delete Report**.


Tip: You want to save a scheduled report to your file system or send it by email report now, but you don't want to change the original schedule? Click **Reports**. Select the report, and click **Actions** > **Run Now**. The report is sent without changing the schedule that you saved.

Tutorial: Creating a predefined capacity report about storage systems

Create, schedule, and send predefined capacity reports about your storage systems by email, to keep your colleagues informed about capacity usage and to help them plan capacity and prevent capacity shortages.


You can create a predefined report about the capacity usage of all of the storage systems in your storage environment, a selection of storage systems in your storage environment, or the storage systems that are assigned to a group such as a general group.

- 1. Click Reports, and then click Create Report.
- 2. In the Capacity Reports section, click **Block Storage Systems**.
- 3. Enter the unique name of the report.
- 4. To specify the range of the report, choose one of the following options:



Option	Description
Capacity information about all of the storage systems	Generate capacity information about all of the block, file, and object storage systems in your storage environment.
Capacity information about one or more storage systems	Generate information about a set of storage systems.
Capacity information about the storage systems that are assigned to a group	Generate information about the storage systems that were assigned to a group such as the production, preproduction, or test platforms general group.

- 5. Add and remove columns to choose the capacity and any other information that you want to include in the report.
- 6. Add up to four filters to deliver the capacity information that the report's recipients need.



Tip: To make sure that the report's recipients get the information that you want them to get, use the report preview feature. Each time you apply a filter, or make a configuration change to the report, such as add or remove a column or filter, the table view of the report is refreshed.

7. Choose one of the following options:

Option	Description
Email	Save and send the email now or attach the report to the email as a CSV, PDF, or HTML file.
Save to File System	Schedule and save the report to the default reports folder or to the folder that you created for saving reports.
Email and Save to File System	Schedule the report, send it by email or attach the report to the email as a CSV, PDF, or HTML file, and save the report to the default reports folder or to the folder that you created for saving reports.

The **Reports** page opens where you can preview and edit the report that you created.

Tutorial: Creating a predefined capacity report about pools

Create, schedule, and send predefined capacity reports about your pools by email, to keep your colleagues informed about capacity usage and to help them plan capacity purchases and prevent capacity shortages.

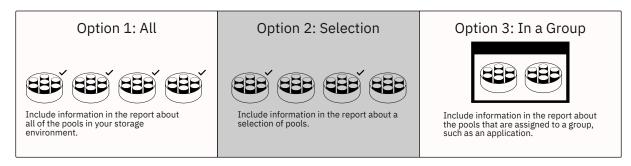
You want to send a regular report to the owners of block storage resources so that they'll know:

- · Which pools are depleted
- When the pools will run out of capacity

To do this, you include the **Zero Capacity** column and create a filter.

What's zero capacity: The capacity information that is collected over 180 days is analyzed to determine, based on historical storage consumption, when the pools will run out of capacity. The pools that have already run out of capacity are marked as depleted. For the other pools, a date is provided so that you know when the pools are projected to run out of capacity. If sufficient information isn't collected to analyze the storage usage of the pool, None is shown as the value for zero capacity. If a capacity limit is set for the pool, the date shown in the Zero Capacity column is the date when the available capacity based on the capacity limit will be depleted. For example, if the capacity limit for a 100 GiB pool is 80%, it is the date when the available capacity of the pool is less than 20 GiB. Depleted is shown in the column when the capacity limit is reached.


- 1. Click Reports, and then click Create Report.
- 2. In the Capacity Reports section, click **Pools**.
- 3. Enter the unique name of the report.
- 4. To specify the range of the report, choose one of the following options:



Option	Description
Capacity information about all of the pools	Generate capacity information about all of the pools in your storage environment.
Capacity information about one or more pools	Generate information about a set of pools. You can select the pools by the storage systems that they belong to or by the tiers that the pools are assigned to.

Option	Description
Capacity information about the pools that are assigned to a group	Generate information about the pools that were assigned to a general group such as an application.

- 5. Add and remove columns to choose the capacity and any other information that you want to include in the report.
- 6. Add up to four filters to deliver the capacity information that the report's recipients need.



Tip: To make sure that the report's recipients get the information that you want them to get, use the report preview feature. Each time you apply a filter, or make a configuration change to the report, such as add or remove a column or filter, the table view of the report is refreshed.

7. Choose one of the following scheduling options:

Option	Description
Email	Save and send the email now. You can also attach the report to the email as a CSV, or HTML file, or both.
Save to File System	Schedule and save the report to the default reports folder or to the folder that you created for saving reports.
Email and Save to File System	Schedule the report, send it by email or attach the report to the email as a CSV file, PDF file, or HTML file, and save the report to the default reports folder or to the folder that you created for saving reports.

The **Reports** page opens where you can preview and edit the report that you created.

Tutorial: Creating a predefined capacity report about the volumes assigned to servers

Create, schedule, and send predefined reports about the capacity of the volumes that are assigned to servers.

When you create reports about the capacity of volumes, you can select volumes:

- By the storage systems that they belong to
- · By the pools that they belong to
- · By the servers that they are assigned to
- By the applications that they are assigned to

In this tutorial, you create a capacity report about the volumes that are assigned to the servers in your storage environment.

- 1. Click **Reports**, and then click **Create Report**.
- 2. Click the **Volumes** capacity report.
- 3. Enter the name of the report.
- 4. Select Capacity information about one or more volumes.
- 5. On the **Select Volumes** page, select **Servers** from the **Volumes** drop-down list.
- 6. On the **Select Volumes by Servers** page, you can:
 - a) Choose all of the servers.
 - b) Choose one or more of the servers.

- c) Select **Name** from the **Filter** drop-down list. Then, enter a pattern that matches the names of the servers that you want to select.
- 7. Choose the columns that you want to include in the report.

You can rearrange the columns by dropping the columns where you want them to appear in the report.

8. Schedule the report.

Tutorial: Creating a predefined capacity report about managed disks by storage systems

Share information with your colleagues about the capacity of your managed disks grouped by the storage systems that they belong to.

- 1. Click Reports, and then click Create Report.
- 2. Pick the Managed Disks capacity report.
- 3. Enter the name of the report and choose Capacity information about one or more managed disks.
- 4. On the **Select Managed Disks** page, change the selection from **Managed Disks** to **Block Storage Systems**.
- 5. On the **Select Managed Disks by Block Storage Systems** page, you can:
 - a) Choose all of the storage systems that have managed disks.
 - b) Choose one or more of the storage systems that have managed disks.
 - c) Select **Name** from the the column filter. Then, enter a pattern that matches the names of the storage systems that you want to select.
- 6. Choose the columns that you want to include in the report.

You can rearrange the order of the columns by dropping the columns where you want them to appear in the report.

- 7. Add and apply filters.
- 8. Schedule the report.

Tutorial: Creating a predefined capacity report about tiered pools

Create a pools by tier report to share information about the capacity of the pools that are assigned to your production platform.

The pools that are used on you production platform are assigned to one or more tiers.

Assign pools to tiers: Click Storage > Pools. To add the Tier column to the table, right-click any column heading, and select Tier. Select and then right-click the pools on your production platform, and click Set Tier. If you used a naming pattern for the pools or storage systems that are assigned to your production

platform, you can filter the pools by pool or storage system name. Just click . Then select **Name** or **Storage System** and enter the pattern that matches the names of the pools or storage systems in the filter

You want to create a scheduled report that shows the key capacity values of the pools on your production platform.

- 1. Click **Reports**, and then click **Create Report**.
- 2. Pick the **Pools** capacity report.
- 3. Enter the name of the report, and click Capacity information about one or more pools.
- 4. Select **Tiers** from the **Pools** drop-down list.
 - A list of the tiers that were created is shown.
- 5. Select the tiers that you created for your production platform.
 - When you click **Next**, capacity information about the pools assigned to the tiers that you selected is shown.
- 6. Select the capacity information that you want to include in the report.

Tip: Include the name of the tier in the report. In the General section, click **Tier**. Drag the **Tier** column and drop it before the **Name** column. To sort the pools by the tiers they are assigned to, click the **Tier** column.

- 7. To refine the information that you want to include in the report, add filters.
- 8. Schedule the delivery of the report.

"Optimizing storage tiering" on page 438

To optimize the placement of volumes on storage tiers, analyze the tiering of volumes in your storage environment.

"Setting the tier level of storage pools" on page 441

To analyze tiering and to balance pools, you must set the tier level of storage pools.

Creating predefined inventory reports

To share information with your colleagues about the configuration and properties of your storage assets, create inventory reports.

Before you start

If you want to send reports by email, you must set up the mail server. If you already set up the mail server to send alerts or reports, or if you want to only save reports to your file system, you don't have to set up the mail server.

If you want the reports that you will create to be saved to your file system to be stored in a folder other than the default reports folder, you must create the folder before you create the reports.

Saving reports to file systems on remote servers: The machine on which the Web server component was installed and runs must be able to access the remote or shared file system that is specified for saving reports. To make file systems on a server available to remote clients, they must be mounted as a CIFS export on Windows operating systems or as an NFS export on UNIX or Linux operating systems. On Windows operating systems, the file system must be mounted as the local system user or from a service that is running under the local system account so that it is visible to all logon sessions. One way of mapping a drive as the local system user is to map the drive from a psexec command window. The psexec command window is available as part of the Sysinternals command-line tools for Windows operating systems.

Step 1: Pick your predefined report

Click **Reports**, and then click **Create Report** and pick the inventory report that you want to create.

Step 2: Name the report

Provide a unique name for your report. You can use alphabetical and numerical characters, hyphens, dashes, and blank spaces.

Step 3: Specify the scope of the report

Generate information about all of the resources of a specific type, such as an inventory report for block storage systems, or IBM Spectrum Virtualize nodes, or IBM Spectrum Scale nodes.


Step 4: Select columns

To ensure that the report's recipients get the information that they need, you can add more information, such as the location of the resources or custom tags. You can also remove the information that your colleagues don't need.

Step 5: Add filters

To refine the information in the report, you can add up to four filters.

Sample filters for block storage systems inventory reports



For example, if you want to check whether your IBM Storwize family block storage systems have the correct firmware version, you can add filters to list the storage systems with firmware versions that don't match the value that you enter.

Don't forget to check the table when you apply filters to make sure that the recipients of the report will get the information that they need.

Step 6: Schedule and deliver the report

To notify your colleagues of inventory issues, create one-off inventory reports or create a schedule to provide regular updates about the state of your storage environment.

Tip: Instead of sending the report by email, you can save the report to the default reports folder or to a folder on your file system. Alternatively, when you send the report by email, you can also attach the report as a CSV, PDF, or HTML file.

Review, edit, and delete reports

To review the report that you created, click **Reports**, expand the relevant section, such as **Inventory Report**, and then select the report.


You can edit the report, such as change the name and scope of the report or schedule of the report. To delete the report, click **Actions** > **Delete Report**.

Tip: You want to save a scheduled report to your file system or send it by email report now, but you don't want to change the original schedule? Click **Reports**. Select the report, and click **Actions** > **Run Now**. The report is sent without changing the schedule that you saved.

Tutorial: Creating an inventory report about block storage systems

Create a predefined report that includes information such as the name, type, model, vendor, and location of your block storage systems.

- 1. Click **Reports**, and then click **Create Report**.
- 2. In the Inventory Reports pane, choose Block Storage Systems.
- 3. Enter the unique name of the report.
- 4. Choose the inventory information that you want to include in the report.
- 5. Add up to four filters.



Tip: Don't forget to click **Apply Filter** when you add each filter, and check the report preview to ensure that the report's recipients will get the information that they need.

6. Choose one of the following scheduling options:

Option	Description
Email	Save and send the email now. You can also attach the report to the email as a CSV, PDF, or HTML file.
Save to File System	Schedule and save the report to the default reports folder or to the folder that you created for saving reports.
Email and Save to File System	Schedule the report, send it by email or attach the report to the email as a CSV, PDF, or HTML file, and save the report to the default reports folder or to the folder that you created for saving reports.

The **Reports** page opens where you can preview and edit the report that you created.

Creating custom reports

Create, configure, and save or send custom reports by email that include asset, capacity, configuration, health status, or performance information about your storage resources.

If you want to send custom reports by email, you must set up the mail server. If you already set up the mail server to send alerts, or chargeback and consumer reports, or if you want to only save reports to your file system, you don't have to set up the mail server.

If you want the reports that you will create to be saved to your file system to be stored in a folder other than the default reports folder, you must create the folder before you create the reports.

Saving custom reports to file systems on remote servers: The machine on which the Web server component was installed and runs must be able to access the remote or shared file system that is specified for saving custom reports. To make file systems on a server available to remote clients, they must be mounted as a CIFS export on Windows operating systems or as an NFS export on UNIX or Linux operating systems. On Windows operating systems, the file system must be mounted as the local system user or from a service that is running under the local system account so that it is visible to all logon sessions. One way of mapping a drive as the local system user is to map the drive from a psexec command window. The psexec command window is available as part of the Sysinternals command-line tools for Windows operating systems.

Use custom reports to quickly create reports to alert members of your organization about performance anomalies, health issues, capacity shortfalls, or capacity depletion. You can also schedule custom reports to keep a close watch on critical resources such as applications, or the storage resources that you use in your production environment.

From any table view that shows information about your storage systems or their related resources, you can create, configure, schedule, and save or send reports by email.

For example, you can create custom reports for top-level storage resources such as:

- · Storage systems
- · Hypervisors
- Servers
- Fabrics
- Switches
- Back-end storage systems

You can create custom reports for internal resources such as:

- Disks
- Pools

- Volumes
- Filesets
- · File systems
- · Network Shared Disks

And, if you added groups, you can create custom reports for general groups, applications, and departments.

File names for reports: The format for the name of reports that are saved to the file system is: <report name>_<epoch time>.<HTML|PDF|CSV>, where epoch time reflects the time of the server where IBM Spectrum Control is installed. For example, the name of your report is MyReport, and you save it as a CSV file on 16 October, 2018 at 08:15pm (UTC) server time. The name of the file for the report is MyReport_1539720900000.CSV.

Tip: Monitor the default reports folder or the custom folders that you use to save reports to ensure that you have enough space for saving reports. For example, if you schedule and save a daily report to your file system, each version of the report is saved to the folder you choose. The factors that determine space usage are:


- The number of rows in the report.
- The frequency of the report.
- The format of the report.

Try it out: Open any page that shows information about your storage resources, such as the **Block Storage Systems** page, click **Create Report**, and create a custom report. Or, switch to the **Performance** page and click **Create Report**, to create a performance report.

- 1. Open any page that shows capacity or performance information about your storage systems, and then click **Create Report**.
- 2. Provide a unique name for your report. You can use alphabetical and numerical characters, hyphens, dashes, and blank spaces.

Remember: When you create your report, you can also use the table view features to customize your report, such as:

- Drag the column headings to reorder the information that you want to show in the report.
- Click the column heading in the table view to sort the values in the column.
- 3. To ensure that the report's recipients get the information that they need, you can add other information, such as the location of the resources or custom tags and remove the information they don't need.
- 4. To refine the capacity or asset information that you want to share, you can add up to four filters.



For example, if you want to notify resource administrators or owners about capacity shortfalls, you can set threshold values for capacity information. Or if you want to exclude information that is not of interest to the report's recipients, you can add a filter. Check the report preview when you apply filters to make sure that the recipients of the report will get the information that they need.

Tip: You can add advanced filters to refine information about capacity and information about the assets in your storage environment. If advanced filters aren't available for a resource, you can use the

filter feature that is available in the table. For example, in reports about volumes, you can filter by volume, pool, or storage system name, or by Easy Tier attributes.

5. Choose one of the following scheduling options:

Option	Description
Email	Save and send the email now. You can also attach the report to the email as a CSV, PDF, or HTML file.
Save to File System	Schedule and save the report to the default reports folder or to the folder that you created for saving reports.
Email and Save to File System	Schedule the report, send it by email or attach the report to the email as a CSV, PDF, or HTML file, and save the report to the default reports folder or to the folder that you created for saving reports.

To review the report that you created, click **Reports** > **Reports**, expand **Customs Report**, and then select the report.

You can edit the report, such as change the name and the schedule for the report. To delete the report, click **Actions** > **Delete Report**.

Tip: You want to save a scheduled report to your file system or send it by email report now, but you don't want to change the original schedule? Click **Reports** > **Reports**. Select the report, and click **Actions** > **Run Now**. The report is sent without changing the original schedule.

Related concepts

"Monitoring the capacity of resources" on page 348

IBM Spectrum Control can collect information about the capacity and space usage of block, file, and object storage resources. This information includes key metrics that can help you measure, identify, and troubleshoot capacity and space usage issues in your storage. You can view capacity metrics for storage resources, including tiers.


Related tasks

"Configuring the email server" on page 499

To send reports, you must configure the email server.

Tutorials: Creating custom capacity and performance reports for applications

In these tutorials, you create an application so that you can create, schedule, and send reports by email about the storage resources that the application consumes.



You create the application IBM Spectrum Control, and then create custom capacity and performance reports.

Creating the application for the capacity report

In this scenario, you create the application that you want to use to generate the custom report.

Add the storage resources that you application consumes to the application that you create in IBM Spectrum Control.

- 1. From the **Groups** menu, click **Applications**.
- 2. Click Create Application.
- 3. Complete these actions:
 - Enter the name.
 - · Add a description.

- Enter the type. For example, to distinguish this application from the other types of application that you create, enter DB or the type of DB, such as db2.
- Enter the subtype. For example, you can enter preproduction, test, or, production to categorize applications by environment.
- 4. Click **Assign storage resources to the application**, and then click **A selection of volumes that match** a a pattern.
- 5. In this scenario, you enter an asterisk * as the name pattern for the volumes, then you click **Belonging** to, and add a name pattern that matches the names of the pools that contain the volumes.

Assigning resources: Enter name patterns instead of adding comma-separated lists of the resources that an application uses. When you add more storage resources that match the name pattern, they are dynamically added to the application.

6. Click **Preview** to check that you added the storage resources that the application consumes, and then click **Save**.

Creating the capacity report for the application

In this scenario, you create a scheduled report about the capacity and storage usage of an application's pools.

You want to create and schedule a report about the capacity of the pools that are assigned to an application.


- 1. From the **Groups** menu, click **Applications**.
- 2. In the **Related Resources** section, double-click the application, and then click **Pools**.
- 3. Click Create Report, enter the unique name of the report, and then click Next.
- 4. Select the capacity values that you want to include in the report.

In this scenario, the following actions are completed to configure the report output:

- a) In the **General** section, exclude all of the columns except for the **Name** and **Tier** columns.
- b) In the **Storage** section, add capacity values, and select **Zero Capacity** from the **Other** section.

What's zero capacity: The capacity information that is collected over 180 days is analyzed to determine, based on historical storage consumption, when the pools will run out of capacity. The pools that have already run out of capacity are marked as depleted. For the other pools, a date is provided so that you know when the pools are projected to run out of capacity. If sufficient information isn't collected to analyze the storage usage of the pool, None is shown as the value for zero capacity. If a capacity limit is set for the pool, the date shown in the Zero Capacity column is the date when the available capacity based on the capacity limit will be depleted. For example, if the capacity limit for a 100 GiB pool is 80%, it is the date when the available capacity of the pool is less than 20 GiB. Depleted is shown in the column when the capacity limit is reached.

- c) To reorder the report output, drag the column headings.
- 5. Add up to four filters.



Tip: Don't forget to click **Apply Filter** when you add each filter, and check the report preview to ensure that the report's recipients will get the information that they need.

6. Choose one of the following scheduling options:

Option	Description
Email	Save and send the email now. You can also attach the report to the email as a CSV, PDF, or HTML file.
Save to File System	Schedule and save the report to the default reports folder or to the folder that you created for saving reports.
Email and Save to File System	Schedule the report, send it by email or attach the report to the email as a CSV, PDF, or HTML file, and save the report to the default reports folder or to the folder that you created for saving reports.

7. Click Save and Schedule.

Depending on the scheduling option that you chose, the report is saved to your file system, sent, or will be sent at the interval that you specified. To preview or edit the report, click **Reports** > **Reports**, expand the **Custom Report** section, and select the report.

Creating the performance report for the application

In this scenario, you create a performance report about the volumes that are assigned to the application.

- 1. Click **Groups > Applications**.
- 2. Double-click the application, and then click **Volumes** in the navigation pane.
- 3. Select the volumes and click **View Performance**.

Tip: Before you create the performance report, add more performance metrics to the table. Click **Edit Table Metrics** and select the additional metrics that you want in the report.

- 4. Click Create Report.
- 5. Enter the name of the report and click **Next**.
- 6. Select the metrics that you want to include in the report and click **Next**.
- 7. Choose one of the following scheduling options:

Option	Description
Email	Save and send the email now. You can also attach the report to the email as a CSV, PDF, or HTML file.
Save to File System	Schedule and save the report to the default reports folder or to the folder that you created for saving reports.
Email and Save to File System	Schedule the report, send it by email or attach the report to the email as a CSV, PDF, or HTML file, and save the report to the default reports folder or to the folder that you created for saving reports.

The report is created. To edit the report, click **Reports** > **Reports**, expand the **Custom Report** section, and click the report.

Custom capacity and performance view reports


When you create custom capacity or performance view reports, you can specify a relative time range for the information that is shown about the resources in the reports.

To generate custom reports with relative time ranges, you must select one or more resources and click **View Capacity**, or **View Performance**. Alternatively, you can open the page for the resource and click the **Capacity** or **Performance** tab.

Remember: If you decide to specify your own start date, or end date, or both, you'll get the same information in the report each time it is generated. For example, you change the time range to 2 days, January 1 to January 2, and then generate a weekly report. Each time the report is generated, the information that was collected from January 1 to January 2 is used to generate the report.

Relative time ranges for custom capacity reports


When you create custom capacity reports on **Capacity View** pages, the default time range for the capacity information that is shown in the report comprises an aggregate of the capacity information that was collected over the previous month.



If you navigate to the capacity view for a storage resource, you can specify a relative time range for the capacity information in the report, such as last week, month, 6 months, or year. Depending on the time range that you specify, the aggregated values for the capacity information are shown in the report.

Relative time ranges for custom performance reports

When you create custom performance reports on **Performance View** pages, the default time range for the performance information that is shown in the report comprises an aggregate of the performance information that was collected over the last 12 hours.

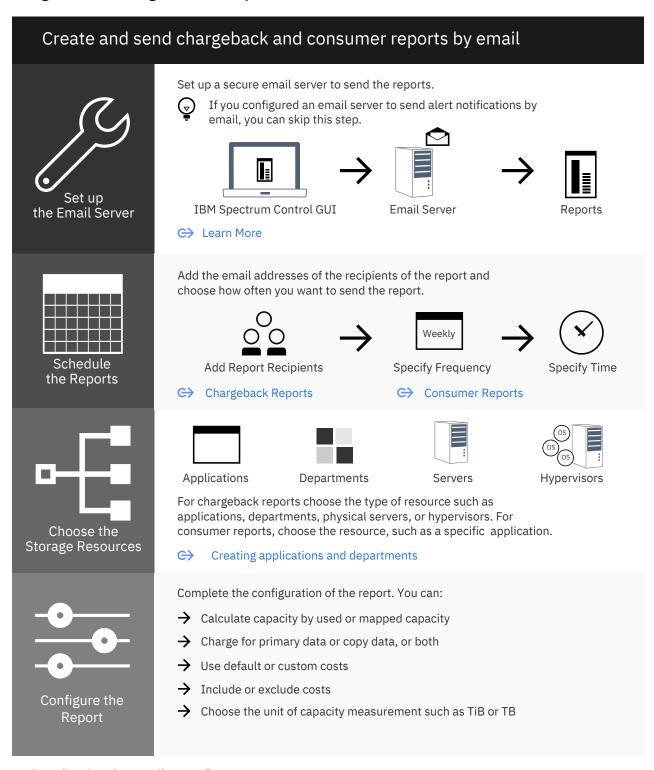


If you navigate to the performance view for a storage resource, you can specify a relative time range for the capacity information in the report, such as the last hour, 6 hours, 12 hours, day, week, month, 6 months, or year. Depending on the time range that you specify, the aggregated values for the performance information are shown in the report.

Creating chargeback and consumer reports

To help plan capacity purchases and make your organization aware of the cost and the amount of the storage that is used by storage consumers, create chargeback and consumer reports.

Chargeback and storage consumer reports



- 1. "Configuring the email server" on page 499
- 2. "Creating chargeback reports" on page 494

- 3. "Creating consumer reports" on page 496
- 4. "Monitoring and administering applications, departments, and general groups" on page 391

In chargeback reports, you create reports that show the capacity and the cost of the block and file storage that is used by the types of storage consumers in your storage environment such as:

- Applications
- Departments
- Hypervisors
- · Physical servers

In storage consumer reports, you create reports that show the capacity and cost of the block storage that is used by a single storage consumer, such as the capacity and the cost of the storage that is used by a single application.

To make the owners of the storage consumers aware of the capacity and cost of the storage they use, you schedule and send chargeback and storage consumer reports by email.

Show cost and capacity of tiered storage: If you want to show the block capacity of tiered storage and differentiate between the cost of storing data on tier 1 or on lower tiers of storage, you must assign your block storage pools to tiers. Click **Storage** > **Pools**. Right-click one or more pools, and then click **Tier** and the tier level from the menu.

To view the capacity and charge for the storage that is used by applications and departments, you must create models of the applications and departments in your business organization and map the storage resources to the applications and map the applications that the departments use to the departments. If you organize an application or department in hierarchies, the chargeback and consumer reports show the capacity and cost of the storage resources that the application and its child applications use and the capacity and cost of the storage resources that the department and its child departments use.

Storage capacity

You can choose the unit of measurement that is used to show the amount of storage that is consumed by your storage resources. In consumer and chargeback reports, you can show the storage that is consumed, for example, in PB, PIB, TB, TiB, GB, or GiB.

Storage costs

The cost of the storage that is used by storage consumers is included in all reports even if the storage costs are set to zero. If you want to send reports that show the capacity of the storage, but not the cost of the storage that is used, you can exclude the storage costs from the report.

Before you decide how you want to set the costs for the block storage that is used by the storage consumers, you must choose whether you want to charge for the storage that is allocated to the storage consumer or the storage that is assigned to the storage consumer.

By default, block storage is calculated by the amount of storage that is allocated to the volumes that are being used for storage by the consumer. Whether you choose to calculate by allocated or assigned capacity, it doesn't affect the calculation of capacity for standard-provisioned volumes because standard-provisioned volumes are fully allocated when they are created. However, the calculation of the capacity of thin provisioned and compressed volumes is affected depending on whether you choose to calculate capacity by allocated or assigned capacity.

For example, an application has three volumes, two fully allocated volumes with 1 TiB each and one thin-provisioned volume with an assigned capacity of 3 TiB. The thin-provisioned volume is allocated capacity in increments of 1 TiB. If you generate a report that calculates capacity by allocated capacity, and the thin-provisioned volume is currently allocated 1 TiB, the total amount of capacity for the application is 1 TiB + 1 Tib + 3 TiB. However, if you choose to calculate by assigned space, the total capacity for the application is 1 TiB + 1 Tib + 3 TiB = 5 TiB. The cost is also affected because the cost is the total capacity of the block storage, which is multiplied by the unit cost or the unit cost per tier.

Restriction: If you choose to calculate block storage by used capacity, the block storage for FlashSystem A9000 and FlashSystem A9000R storage systems is calculated based on the provisioned capacity percentage of the volumes. This restriction occurs because used capacity information for volumes is not available from FlashSystem A9000 and FlashSystem A9000R.

Default or custom storage costs

When you create or update a report, you can choose one of the following options:

Use default cost

Set or use default costs per unit and type of storage.

If you choose to use the default costs option, the changes that you make apply to all of the reports that use the default option to calculate the cost of storage. When you exclude costs from reports that use default costs, you exclude costs from all reports that use default costs to calculate the cost of storage.

Set custom cost

Use custom costs per unit and type of storage.

If you choose to use the custom costs option, the costs that you set are used to calculate the cost of storage only for that report. If you exclude costs from the custom report, the costs are only excluded from the custom report that you create or edit.

The reports that were created in IBM Spectrum Control Version 5.2.11 use the unit costs that were specified when the reports were configured.

Tip: If you want to change the configuration of a report, such as how storage costs are calculated, click **Reports** > **Email Reports**, select the report, and then click **Edit**.

Deleted and unavailable reports

If you delete a storage consumer, such as an application, the entry for that storage consumer is removed from all chargeback reports for applications. If you created a consumer report for a storage consumer that was deleted, the report is automatically deleted.

If you organize your applications or departments in hierarchies, the chargeback and consumer reports show the total capacity and costs for the parent application or department and its child applications and departments. However, if you create a consumer report for a specific application or department and then designate the application or department as a child application of another application, the consumer report is deleted.

For example, you create a consumer report for an application called German Online Sales. You then reorganize your sales applications in a hierarchy and the application for German Online Sales becomes a child application of European Sales. To include the capacity and cost of the storage that is used for German Online sales, you must create a consumer report for the application for European Sales.

Creating chargeback reports

Create chargeback reports that show how much capacity is used and the cost of the capacity that is used by applications, departments, hypervisors, and physical servers.

Chargeback reports help managers to realize the cost of the storage that they use, plan storage usage more efficiently, and reduce storage costs. For example, a manager gets a chargeback report that shows that the bulk of the storage that is used by a non-critical application is on tier-1. To reduce storage costs and use storage more efficiently, the manager can request the storage administrator to use tier-2 or lower tiers of storage for the application's data.

- 1. Click Reports, and then click Create Report.
- 2. In the Summary Reports pane, click Chargeback.
- 3. Type the unique name of the report, and enter the delivery and the scheduling details.
- 4. Choose the type of storage consumer from the list.

If you choose **Hypervisor** or **Physical Server**, the capacity and cost of the hypervisors and physical servers that belong to clusters are listed under the name of the cluster.

- 5. Choose how you want to calculate the capacity and cost of thin-provisioned storage.
- 6. Choose the unit of capacity measurement that you want to use in the report.

Tip: To make sure that the recipients of the report get the report that they want, check the report preview.

7. Choose whether you want to use default or customized costs for block and file storage.

To change the default costs for block and file storage, click the lock (🛍) icon.

8. By default, the storage costs are shown in the report. To exclude storage costs from the report, choose one of the following options:

Storage cost option	Steps
Use Default Cost	a. Click the lock () icon. b. In the Include Storage Cost in Report section, click No .
	Note: If you exclude costs from a report that uses default costs, the costs are excluded from all of the chargeback and consumer reports that use default costs.
Use Custom Cost	In the Include Storage Cost in Report section, click No .

- 9. Choose the types of storage that you want to include in the report.
- 10. Set the unit costs for the storage that is used.

To get a complete picture of the capacity and cost of the block storage, include the cost of the data that is copied.

11. Click Save.

If you set the frequency of the report to **Now**, you click **Save and Send**.

The report is saved and, depending on the scheduling option, the report is sent to the recipients or scheduled to be sent later. If you want to change the configuration of the report, click **Reports**, select the report, and then click **Edit**.

Tip: You want to send a scheduled report now, but you don't want to change the original schedule? Click **Reports** > **Reports**. Select the report, and click **Actions** > **Run Now**. The report is sent without changing the original schedule.

If you want to extract capacity information regularly about your storage resources, you can use the Representational State Transfer (REST) API to automatically export the data that you need to generate custom reports or share with external applications.

Related information

REST API

Chargeback reports

Configure the chargeback reports to include block and file storage and to calculate the cost of the block and file storage that is used by the storage consumers.

You configure, schedule, and send chargeback reports so that the owners and managers of storage resources:

- Know the cost of maintaining the block data and file data that they use
- Know the cost of maintaining block data on each tier of storage
- Know the cost of maintaining block copy data

Include total block and file capacity in the report

In the **Configure Report** section, you choose the types of storage that you want to show in the report. For example, if you click the **File Storage** check box, the names of the resources and the total file capacity that the resources use are shown in the report.

For block storage, you can show the total storage capacity and, if your storage is tiered, the total capacity by tier. You can also break down the total block capacity into total primary capacity and total copy capacity.

Tip: The recipient of the report, such as the application owner, sees that the capacity for copy data is high for tier-1 storage. To reduce costs and use the block storage more efficiently, the application owner can ask the storage administrator to place the volumes that are used for copy data on lower and less expensive storage tiers. The storage administrator can also analyze the tiering of the pools to make sure that the application's performance is not affected when the copy data volumes are moved to lower storage tiers.

Include block and file costs in the report

In the **Set Storage Cost** section, you set the cost per TiB of storage that the resources use. The total cost of the storage that is used is shown in the Cost column in the report.

For block storage, you set the costs as follows:

- · If your storage environment is tiered, you set the price per TiB for each tier of storage
- If your storage environment is not tiered, you set the price per TiB for all of your block storage
- If your storage environment contains tiered and non-tiered storage, you can set the price per TiB for the tiered storage, but you can't set a price for the non-tiered storage. Although you can't set the cost for the non-tiered storage, the capacity of the non-tiered storage that is used is calculated.

If you assign your block storage pools to tiers, you can differentiate between the cost of expensive tier-1 storage and the cost of the less expensive storage that is used on the lower tiers of storage.

For file storage, you set a price for all of the file storage that is used by your resources.

When you configure chargeback reports, you can send the report immediately or schedule the report to be sent every day, every week, or every month.

Tip: If your storage environment changes, you should check whether you need to update your scheduled reports. For example, if you add a storage tier, you can edit the report and set the price for the new tier.

Creating consumer reports

Create consumer reports that show how much block capacity is used and the cost of the block capacity that is used by an application, department, hypervisor, or physical server.

Storage consumer reports help the managers and owners of applications, departments, hypervisors, and physical servers realize the cost of the block storage that they use, plan block storage usage more efficiently, and reduce block storage costs.

- 1. Click Reports, and then click Create Report.
- 2. In the Summary Reports pane, click **Storage Consumer**.
- 3. Type the unique name of the report, and enter the delivery and the scheduling details.
- 4. From the list, click the type of storage consumer and type or click the name of the storage consumer. To generate a report about the hypervisors or physical servers that belong to a cluster, you choose **Hypervisor** or **Physical Server**, and then choose or type the name of the cluster.
- 5. Choose whether you want to calculate block capacity by allocated or assigned space.
- 6. Choose the unit of capacity measurement that you want to use in the report.

Tip: To make sure that the recipients of the report get the report that they want, check the report preview.

7. Choose whether you want to use default or customized costs for block storage.

8. Choose the types of storage that you want to include in the report.

To change the default costs for block storage, click the lock (🖹) icon.

9. By default, the storage costs are shown in the report. To exclude storage costs from the report, choose one of the following options:

Storage cost option	Steps
Use Default Cost	 a. Click the lock () icon. b. In the Include Storage Cost in Report section, click No. Note: If you exclude costs from a report that uses default costs, the costs are excluded from all of the chargeback and consumer reports that use default costs.
Use Custom Cost	In the Include Storage Cost in Report section, click No.

10. Set the unit costs for the block storage that is used.

To get a complete picture of the capacity and cost of the block storage, include the cost of the data that is copied.

11. Click Save.

If you set the frequency of the report to **Now**, you click **Save and Send**.

The report is saved and, depending on the scheduling option, the report is sent to the recipients or scheduled to be sent later. If you want to change the configuration of the report, click **Reports**, select the report, and then click **Edit**.

Tip: You want to send a scheduled report now, but you don't want to change the original schedule? Click **Reports** > **Reports**. Select the report, and click **Actions** > **Run Now**. The report is sent without changing the original schedule.

If you want to extract capacity information regularly about your storage resources, you can use the Representational State Transfer (REST) API to automatically export the data that you need to generate custom reports or share with external applications.

Related information

REST API

Consumer reports

Configure consumer reports to calculate the cost and capacity of the block storage that is used by a storage consumer such as an application, a department, a hypervisor, or a physical server. Configure consumer reports to calculate the cost and capacity of the block storage that is used by a storage consumer such as an application, a department, or a physical server.

You configure, schedule, and send consumer reports so that the owner or manager of the storage consumer knows the cost of maintaining the block storage that the storage consumer uses.

In the **Configure Report** section, you include the total block storage capacity and, if your storage is tiered, the total capacity by tier in your report. You can also include the total primary capacity and the total copy capacity of the storage that is used, and you can include the total capacity:

- For VDisk mirrors
- For FlashCopy volumes
- For remote mirrors such as Global Mirror, Global Copy, and Metro Mirror

In the **Set Storage Cost** section, you set the cost per TiB of the block storage that the storage consumer uses.

You set the costs as follows:

- If your storage environment is tiered, you set the price per TiB for each tier of storage
- If your storage environment is not tiered, you set the price per TiB for all of your block storage

• If your storage environment contains tiered and non-tiered storage, you can set the price per TiB for the tiered storage, but you can't set a price for the non-tiered storage. Although you can't set the cost for the non-tiered storage, the capacity of the non-tiered storage that is used is calculated.

Creating summary reports of the storage capacity

Configure reports that make managers aware of the capacity of the storage that is used by their applications, departments, hypervisors, or physical servers. The cost of the storage is not shown in the reports.

You can create reports that show the total capacity of the storage that is used by storage consumers and exclude the storage costs from the reports.

- 1. Click Reports > Reports, and then click Create Report.
- 2. Choose one of the following options:

Option	Steps	
Create chargeback reports	In the Summary Reports pane, click Chargeback .	
Create consumer reports	In the Summary Reports pane, click Storage Consumer .	

- 3. Enter the report details and schedule the report.
- 4. Choose one of the following options:

Option	Steps	
Chargeback reports	Choose the type of storage consumer.	
Consumer reports Choose the type of storage consumer and the storage consumer.		

- To calculate the capacity that is assigned to the thin-provisioned volumes, click **Assigned space**.
 By default, the amount of space that is allocated to the thin-provisioned volumes is used to calculate the cost.
- 6. Choose the unit of capacity measurement that you want to use in the report.

Tip: To make sure that the recipients of the report get the report that they want, check the report preview.

- 7. In the **Configure Report** section, choose the capacity values that you want to include in the report.
- 8. By default, the storage costs are shown in the report. To exclude storage costs from the report, choose one of the following options.

If you exclude costs from a report that uses default costs, the costs are excluded from all of the chargeback and consumer reports that use default costs.

Storage cost option	Steps
Use Default Cost	 a. Click the lock () icon. b. In the Include Storage Cost in Report section, click No.
Use Custom Cost	In the Include Storage Cost in Report section, click No .

9. Click Save.

If you set the frequency of the report to **Now**, you click **Save and Send**.

The reports include the capacity values and exclude the costs of the capacity that is used.

Running reports

You can run reports that you created without defining a schedule or changing the schedule that was created for the report.

If you created a report, but didn't schedule it, you can run the report again. You can also run reports that you scheduled without changing the schedule for the report.

- 1. Click **Reports**.
- 2. Click the report.

To run two or more reports, press **Ctrl** or press **Shift** and click the reports.

3. Click Actions > Run Now.

The reports are generated and sent to the reports' recipients.

Editing reports

Change the configuration and scheduling for reports.

You can edit and reschedule custom, predefined capacity and inventory, and consumer and chargeback reports.

- 1. From the **Reports** menu, click **Reports**.
- 2. Expand the section that contains the report.
- 3. Click the report, and then click Edit.

Tip: For custom and predefined capacity and inventory reports, click one of the links in the task pane. By default, the first page that is shown is the **Schedule Delivery** page. For example, to add, change, or remove filters, click **Edit Filter**.

4. Make your changes.

If you schedule the report, the changes are saved and included in the next scheduled report. If you chose send now when you scheduled the report, the changes are saved and the report is sent to the recipients.

Don't forget when you send the report that you will need to reconfigure the schedule if you want to send or save the report at regular intervals.

Deleting reports

Remove the reports that you don't need.

- 1. Click **Reports**.
- 2. Click the report.

To select two or more reports, press Ctrl or press Shift and click the reports.

3. Click Actions > Delete Report.

Configuring the email server

To send reports, you must configure the email server.

The email server that is used to send alert notifications is also used to send reports. If you didn't configure the email server for alert notifications, you must configure the email server before you create reports.

IBM Spectrum Control ensures, when it sends alert notifications and reports by email to your email server, that it complies with the security standards that you configure on your email server. For example, if your email server requires authentication and TLS to send emails, then IBM Spectrum Control uses the

authentication credentials that you provide when you set up the email server, and uses TLS to establish the connection with your email server. Because the email server that you set up is automatically trusted, you don't have any additional management tasks such as importing the server certificate into the keystores that are used by IBM Spectrum Control.

Tip: To change the configuration of the mail server, you click Settings > Notification Settings.

- 1. Type the host name, or IPv4, or IPv6 address of the email server.
- 2. Type the port number that is used by the email server to send alert notifications and reports.
- 3. Provide the user's credentials for authenticating with your email server.
 - If your email server requires authentication to establish the connection, you must enter the user name and password.
- 4. Optional: Type the email address, such as the email address of the administrator, that you want to use for receiving replies to alert notifications and reports.
 - If you don't enter a value, and you entered an email address in the **User Name** field, then replies are sent to that address. Otherwise, replies are sent to no-reply@<Spectrum_Control_server_host_name>.
- 5. Test the connection and then click Save.

You can create reports and send the reports by email.

Investigating issues with reports

You can access log and trace files to investigate the issues that you might have when reports are generated and sent by email or saved to the file system.

In IBM Spectrum Control, you can create, schedule, and send chargeback and storage consumption reports by email, and you can create, schedule, and save custom reports to your file system or send them by email.

To create reports, the following IBM Spectrum Control components and services are used. (The Export server service is used only for custom reports.)

Create, update, and delete reports

Web server, Data server, Device server

Generate and deliver reports

Web server, Data server, Export server, SMTP transmission service

Export server is a new service that converts the data for the report that it gets from the Web server into HTML.

The SMTP transmission service handles all communications, such as the delivery of alert notifications by email, and the sending of chargeback, storage consumption, and custom reports from IBM Spectrum Control to the customers' external SMTP servers.

Setting the trace level for investigating reporting issues

Go to the folder where you installed IBM Spectrum Control and complete these steps:

Tip: For Windows operating systems the default installation directory is C:\Program Files\IBM\TPC. For AIX and Linux operating systems the default installation directory is opt/IBM/TPC.

- 1. On Windows OS, go to installation_dir\IBM\TPC\web\conf. On AIX or Linux OS, go to installation_dir/IBM/TPC/web/conf.
- 2. In a text editor, open the logging.properties file.
- 3. Set the value of the trace.reportdata.level property to trace.reportdata.level=debug.
- 4. Save the logging.properties file.

You have set the trace level for the traceReportData_<number>.log to debug.

Access log, trace, and other information about reports

Table 224. Location of log, trace, and other files

File name	Description and location
ReportData.log	High-level information about the creation and distribution of reports.
	Windows: installation_dir\web\log
	AIX and Linux: installation_dir/web/log
traceReportData_ <number>.log</number>	Detailed information about the creation and distribution of custom reports, and SMTP transmission of custom reports. For example, you can check why a custom report isn't saved to the file system, or sent, or received by email.
	Windows: installation_dir\web\log
	AIX and Linux: installation_dir/web/log
traceWebServer_ <number>.log</number>	Information relating to creating, editing, and scheduling custom reports.
	Windows: installation_dir\web\log
	AIX and Linux: installation_dir/web/log
Console.log	Information about whether the report data service, which runs on the Web server, was started or stopped, errors generated when creating and retrieving custom reports, handshake errors that were generated when custom reports are sent by email.
	Windows:installation_dir\wlp\usr \servers\webServer\logs
	AIX and Linux: installation_dir/wlp/usr/servers/webServer/logs
Messages.log	See the description for the Console.log. Both log files report the same types of errors.
	Windows:installation_dir\wlp\usr \servers\webServer\logs
	AIX and Linux: installation_dir/wlp/usr/servers/webServer/logs
traceDeviceServer_ <number>.log</number>	Detailed information about the SMTP transmission of the chargeback and storage consumption reports sent by email. For example, you can check why a chargeback or storage consumption report wasn't sent.
	Windows: installation_dir\device\log
	AIX and Linux: installation_dir/device/log

Table 224. Location of log, trace, and other files (continued)

File name	Description and location	
Messages.log	Information about whether the Device server was started or stopped, handshake errors that are generated when alert notifications, chargeback, or storage consumption reports are sent.	
	Windows: installation_dir\device\log	
	AIX and Linux: installation_dir/device/log	
traceScheduler_ <number>.log</number>	Information about when the report creation process began and completed for each report.	
	Windows: installation_dir\data\log	
	AIX and Linux: installation_dir/data/log	

Stopping and starting the Export server service

On Windows, the Export server runs as a service called **IBM Spectrum Control - Export Server**, which can be stopped and restarted.

- 1. Click the **Start** menu, type services.msc, and then press Enter.
- 2. Click IBM Spectrum Control Export Server, and stop or start it.

Alternatively, you can run the following batch files or scripts to stop or start the Export server service on Windows or AIX and Linux operating systems:

Operating system	Location	File
Windows	By default, it's C:\Program Files\IBM\TPC\scripts	Start: startTPCExport.bat
	riies (ibm (irc (scripts	Stop: stopTPCExport.bat
AIX and Linux	By default it's/Opt/IBM/TPC/ scripts	Start: startTPCExport.sh
	SCIIPES	Stop: stopTPCExport.sh

Types of predefined capacity and inventory reports

The predefined capacity and inventory reports that you can create for you storage resources are listed.

Check the table to see the predefined reports that you can create for the storage resources in your environment.

Table 225. Predefined capacity and inventory reports		
Resource	Capacity	Inventory
Block Storage Systems	Yes	Yes
File Storage Systems	Yes	Yes
Object Storage Systems	Yes	Yes
Switches	No	Yes
Servers	Yes	Yes
Hypervisors	Yes	Yes
Applications	Yes	No

Table 225. Predefined capacity and inventory reports (continued)		
Resource	Capacity	Inventory
Departments	Yes	No
Managed Disks	Yes	Yes
Controllers	No	Yes
Enclosures	No	Yes
Filesets	Yes	No
NAS File Systems	Yes	No
Modules	No	Yes
Network Shared Disks	Yes	No
Pools	Yes	No
Spectrum Scale Nodes	No	Yes
Spectrum Virtualize Nodes	No	Yes
Storage System FC Ports	No	Yes
Storage System IP Ports	No	Yes
Server Ports	No	Yes
Switch Ports	No	Yes
Volumes	Yes	No

Reports FAQ

Find answers to questions about the reports that you can create in IBM Spectrum Control.

Troubleshooting reports

Why can't I send the report as an attachment?

You're notified that a report can't be sent by email because the report's attachments exceed the maximum size limit. The maximum size limit for attachments is set on your email server.

If you can't change the maximum size limit for attachments, you can edit the report and choose **Save to File System** as the delivery method on the **Schedule Delivery** page.

Alternatively, you can reduce the scope of the report to create smaller and more focused reports. For example, when you create reports about volumes, you can select volumes by the storage systems, pools, or applications that they belong to, or the servers that they are assigned to. You can also refine the report by removing columns and adding filters. And, to reduce the size of attachments, choose CSV as the file type.

Learn more

Why can't I save the report to a custom folder on my file system?

You must have read and write access to the folder and you must create the custom folder before you create the report. Don't forget to enter the full path to the folder, such as C:\Program files\IBM\TPC\myreports or /opt/IBM/TPC/myreports on the **Schedule Delivery** page.

Why can't I see all of the rows in the report in my email?

If the report is very large, the report's recipients might not see all of the rows in the report in the email.

Click **Reports**. Select the report and click **Edit**. On the **Schedule Delivery** page, add the report as a CSV or HTML attachment. To see all of the rows in the report, the reports' recipients can open the attachment.

Creating and sending

Why can't I create a new custom report on the Reports page?

Unlike chargeback, consumer, and predefined reports, you can't create new custom reports on the **Reports** page. To create new custom reports, open any page that shows asset, capacity, configuration, or performance information, such as the **Block Storage Systems** page, and click **Create Report**. On the **Reports** page, you can edit or delete the custom reports that you created.

Learn more

How do I create a custom report about performance?

Go to a page that shows information about block storage resources. Select one or more of the resources, click **View Performance**, and then click **Create Report**.

Learn more

What types of reports can I create?

IBM Storage Insights Pro users can create chargeback, consumer, predefined capacity, predefined inventory, and custom reports. IBM Storage Insights users can create predefined capacity reports for storage systems and pools and predefined inventory reports for storage systems. To create new predefined reports, click **Reports**, and then click **Create Report**.

Learn more

How can I send the reports that I created now?

You want to send a report now although the report, for example, is scheduled to run next week. Or, you created a report that you didn't schedule and want to send again.

Click **Reports**. To select the reports, press **Ctrl** or press **Ctrl + Shift** and click the reports. From the **Actions** menu, click **Run Now**. The report is generated without changing the schedule that you created for the report.

Exporting information about resources

Notify your colleagues about the current state and potential issues with your storage environment by exporting the capacity and performance information that is shown in the GUI.

You can use the information that is shown about your storage resources and their related resources to create reports in HTML or PDF file formats that you can send to your colleagues. Alternatively, you can export the information about your resources to a CSV file, which you can include in a report that you want to create about your storage resources.

You can go to the resource or internal resource page and create reports that contain the following types of information:

Table 226. Types of information in resource reports		
Location	Types of information	
Resource or internal resource page	Status, configuration, and capacity data	
Performance tab for resource or internal resource	Key performance metrics	
Capacity tab for resource or internal resource	Key capacity metrics	

Restriction: You might not be able to view capacity and performance metrics for all of the resources.

For example, you monitor the pools for a SAN Volume Controller and you see that some of the pools are running out of capacity. To notify your colleagues, you export the information about the pools to an HTML file and send it to your colleagues. You want to include information about the input/output performance of

the volumes for your storage systems in a report. To create the report, you export the information that is shown in the performance chart to a CSV file.

If you want to extract configuration, capacity, or performance information regularly about your storage resources, you can use the Representational State Transfer (REST) API to automatically export the data that you need to create your reports.

Related tasks

"Exporting information to a file" on page 452

You can save information about resources, tasks, or alerts to a PDF, CSV, or HTML file. Information that you can export from the GUI includes all the values that are being shown in the columns for a list of resources, tasks, or alerts.

Related information

REST API

Exporting information about capacity shortages in pools

To include information about the capacity of your resources in reports, use the reporting feature in IBM Spectrum Control.

You check the capacity of pools on your SAN Volume Controller storage system and see that the capacity of some of the pools is nearly depleted. So, you decide to notify your colleagues by including information about the pools in a report.

- 1. Click Storage > Block Storage Systems.
- 2. Double-click the storage system and click **Pools** in the navigation pane.
- 3. On the **Pools** page, right-click one of the column headings and click **Zero Capacity**.

Tip: If you want to export information only about the pools that are depleted, click , select **Zero**

Capacity, and type Depleted in the filter field. When a capacity limit is set, the pool is considered to be Depleted when the capacity limit is reached.

4. Click **Actions** and choose one of the options.

You can get information about the capacity of the pools, whether the pools are depleted, and when the pools will become depleted.

Exporting information about the input/output performance of volumes

To include information about the performance of your resources in reports, use the reporting feature in IBM Spectrum Control.

You want to include information about the input/output performance of the volumes in a report.

- 1. Click Storage > Volumes.
- Select all of the volumes in the list and click View Performance.
 By default, the top 10 volumes are ordered by total input/out rate and the reporting period is set to 24 hours.
- 3. Click Actions > Export > Export as CSV.

Information about the performance of the volumes with the highest input/output rates are exported to a CSV file, which you can download and include in a report.

Capacity metrics for reports

To gain insights into storage usage, review the capacity metrics that are collected, analyzed, and shown in reports.

Capacity metrics for chargeback and consumer reports

Learn more about how the capacity and costs for block and file storage is calculated in chargeback and consumer reports.

Chargeback and storage consumer reports

The same block capacity metrics are shown in chargeback and consumer reports. In chargeback reports, the block capacity and costs are shown for all applications, departments, hypervisors, and physical servers. In consumer reports, the block capacity and costs are shown for a specific application, department, hypervisor, or physical server. The capacity and cost of file storage is not included in consumer reports.

If you create reports about physical servers, servers that are virtual machines are excluded from the report.

Tip: To find out which servers are virtual machines, click **Servers** > **Servers**, right-click any column heading and click **Virtual Machine**.

Virtual machine capacity

If you assign the storage resources on a virtual machine to an application or department, the capacity of those storage resources is included in the calculation of capacity and cost for the application or department.

For example, three virtual machines are allocated capacity from a data store, which has a total capacity of 3 TiB. The capacity for the data store comprises:

- One volume on tier 1 with a capacity of 2 TiB
- One volume on tier 2 with a capacity of 1 TiB

The ratio of tier-1 storage to tier-2 storage is 2:1 or 66.66% to 33.33%.

The total capacity of the data store is split across 3 by 1,024 GiB disks and one of the disks is allocated to Virtual_Machine_1. The used capacity on Virtual_Machine_1 is allocated to a database application. So, the total amount of capacity that is allocated to the database application is 1,024 GiB, which comprises 66.66% of tier-1 storage and 33.33% of tier-2 storage. When the report is generated for the database application, the block capacity for tier-1 storage is 1024 GiB*66.66%, which is 682.67 GiB or 0.67 TiB. The block capacity for the tier-2 storage is 1024 GiB*33.33%, which is 341.33 GiB or 0.33 TiB.

If you calculate the capacity and cost of storage by used capacity, the capacity and cost of the storage that is used by the application or department is shown in the chargeback and consumer reports.

Restriction: You might get unexpected results in chargeback and consumer reports if the volumes that are used to provide the capacity on the disks for the virtual machine are thin-provisioned and shared by two or more applications.

Hypervisor and physical server clusters

In chargeback and storage consumer reports, the capacity and cost of hypervisors and physical servers that are members of clusters are included in the capacity and cost of the cluster that they belong to. For example, hypervisor_1 and hypervisor_2 belong to cluster_1. To create the chargeback report about the hypervisors cluster, you choose **Hypervisors** as the consumer, and the capacity and the cost of the two hypervisors are included in the report entry for cluster_1. To create a report for cluster_1 in a storage consumer report, you choose **Hypervisor** as the type of consumer, and cluster_1 as the name of the consumer.

Block storage

The following information is provided about how block capacity is calculated:

Block or Used Capacity

In chargeback reports, the total block capacity for all of the storage resources is called Block capacity, whereas in consumer reports the total block capacity that the storage consumer uses is called Used Capacity.

The total capacity of these volumes is included in the calculation of block capacity:

- The provisioned capacity of the volumes without copies that are allocated to the consumer
- The provisioned capacity of the volumes with copies that are allocated to the consumer such as:
 - The total capacity of VDisk mirror volumes and their copies
 - The total capacity of FlashCopy volumes and their target volumes
 - The total capacity of Global Mirror, Global Copy, and Metro Mirror volumes and their target volumes

For example, an application is assigned all of the block storage volumes on a physical server. In this scenario, the block capacity value includes the capacity of all of the block volumes that are assigned to the physical server. If the assigned volume on the physical server has a VDisk mirror copy or the volume is a FlashCopy source volume, the capacity of the VDisk mirror copy or the capacity of the FlashCopy target volume is also included in the calculation.

Primary capacity

The total capacity of the following types of volumes is included in the calculation of primary capacity:

- The provisioned capacity of the volumes without copies that are allocated to the consumer
- The provisioned capacity of primary VDisk mirror copies that are allocated to the consumer
- The provisioned capacity of the source volumes in FlashCopy relationships that are allocated to the consumer
- The provisioned capacity of the Global Mirror, Global Copy, and Metro Mirror source volumes that are allocated to the consumer

VDisk Mirrors

The total capacity of the VDisk mirror copies for the volumes that are allocated to the consumer.

FlashCopy

The total capacity of the FlashCopy target volumes for the source volumes that are allocated to the consumer.

Remote Mirrors

The total capacity of the Global Mirror, Global Copy, and Metro Mirror target volumes for the source volumes that are allocated to the consumer.

Primary capacity by tier

The primary capacity that is allocated to the consumer for the specified tier.

Copy capacity by tier

The copy capacity that is allocated to the consumer for the specified tier, which is calculated as follows:

- The total capacity of the non-primary VDisk mirror copies
- The total capacity of the target volumes for volumes in FlashCopy relationships
- The total capacity of target volumes for Global Mirror, Global Copy, and Metro Mirror

In tiered environments, the source volume might be on tier 1 and the target volume on tier 3. In this case, the capacity of the source volume is included in the calculation for the total primary capacity for Tier 1 and the capacity of the target volume is included in the calculation for the total copy capacity for tier 3.




Figure 5. Calculating copy capacity

Restriction: For volumes that have multiple volume copies, only the capacity of the first copy and the capacity of the second copy of the volume are included in the calculation of copy capacity.

File storage

File capacity is the total amount of file storage space that is allocated to the consumer.

For example, if a server mounts a fileset export from a file storage system and a quota is defined for the fileset, the file capacity is the hard limit capacity that is defined for the quota.

If a quota is not defined for the fileset, the file capacity is the capacity of the file system that contains the fileset. If multiple fileset exports from the same file system are mounted and the hard limits that are defined for the quota exceed the total file system capacity, then the file capacity is the total capacity of the file system.

If a hypervisor data store is a mounted NFS export, file capacity is the capacity of the file system on which the export is based.

Appendix A. Reference

View reference information that is related to IBM Spectrum Control. Topics include information about alerts, fabrics, commands, configuration and log files, performance metrics, protocols, standards, and accessibility features.

Opening IBM Spectrum Control on Windows operating systems

You can open IBM Spectrum Control CLIs and GUIs and administer IBM Spectrum Control on Windows operating systems.

Opening IBM Spectrum Control GUIs and CLIs

To manage and monitor storage resources, open IBM Spectrum Control GUIs and CLIs.

You can open the following GUIs and CLIs:

- "Opening IBM Spectrum Control GUI" on page 509
- "Opening Db2 Command Window" on page 509
- "Opening IBM Data Studio Administration client on Windows operating systems" on page 510
- "Opening IBM Tivoli Monitoring Services" on page 510

Opening IBM Spectrum Control GUI

1. Choose one of these options:

Option	Description
Windows Server 2012	 a. On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start. b. On the Start page, right-click, and then click the All apps taskbar button.
Windows Server 2012 R2,Windows Server 2016, Windows Server 2019	Click Start > All Programs.

2. Click IBM Spectrum Control > IBM Spectrum Control.

Opening Db2 Command Window

1. Choose one of these options:

Option	Description
Windows Server 2012	a. On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start .
	b. On the Start page, right-click, and then click the All apps taskbar button.
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start > All Programs.

2. Click IBM DB2 > Command Line Tools > Command Window.

Opening IBM Data Studio Administration client on Windows operating systems

1. Choose one of these options:

Option	Description
Windows Server 2012	a. On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start .
	b. On the Start page, right-click, and then click the All apps taskbar button.
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start > All Programs.

2. Click IBM Data Studio > Data Studio Administration Client.

Opening IBM Tivoli Monitoring Services

1. Choose one of these options:

Option	Description	
Windows Server 2012	a. On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start .	
	b. On the Start page, right-click, and then click the All apps taskbar button.	
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start > All Programs.	

2. Click IBM Tivoli Monitoring > IBM Tivoli Monitoring Services.

Accessing administration tools

To manage and maintain IBM Spectrum Control, access the Windows system administration tools.

To complete tasks in IBM Spectrum Control, you must open the following administration and maintenance facilities:

- "Accessing the Control Panel" on page 510
- "Accessing Administrative Tools" on page 511
- "Accessing Windows Services" on page 511
- "Accessing Computer Management" on page 511
- "Accessing Programs and Program Features" on page 511

Accessing the Control Panel

1. Choose one of these options:

Option	Description
	On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start .

Option	Description
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start .

2. Click Control Panel

Accessing Administrative Tools

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start .
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start .

2. Click Administrative Tools

Accessing Windows Services

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start .
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start .

2. Click Administrative Tools > Services

Accessing Computer Management

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start .
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start .

2. Click Administrative Tools > Computer Management

Accessing Programs and Program Features

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start .

Option	Description
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start .

2. Click Control Panel > Programs and Program Features.

Accessing Window Run

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start .
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start .

2. Click Run.

Appendix B. Accessibility features for IBM Spectrum Control

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 Spectrum Control:

- · Keyboard-only operation in the GUI.
- A Knowledge Center that includes the following accessibility features:
 - The Knowledge Center is provided in XHTML 1.0 format, which is viewable in most web browsers.
 With XHTML, you can view documentation according to the display preferences that are set in your browser. XHTML supports screen readers and other assistive technologies.
 - All documentation for IBM Spectrum Control is available in Adobe Portable Document Format (PDF)
 by using the Adobe Acrobat Reader. To access PDFs, go to Printable documentation.
 - All images in the Knowledge Center are provided with alternative text, so that visually impaired users can understand the contents of the images.
- Interfaces that are commonly used by screen readers.

The setting for the automatic-refresh button in the screen reader is toggled to the ON position by default. If you want the screen reader to read the previous text, complete the following steps:

- 1. Navigate to the Accessibility Settings Navigation region by using the arrow keys. The region is located after the IBM Spectrum Control application title.
- 2. Click **Enter** to toggle the automatic-refresh button to the OFF position. (An alert sounds to make you aware that the turn-off automatic-refresh toggle button was pressed.)
- 3. To move backward to the previously read text so that the screen reader can read it again, use the arrow keys. You can move backwards and forwards through the page.
- 4. When you are ready to move on, click **Enter** to toggle the automatic-refresh button to the ON position and to refresh the page. (An alert sounds to make you aware that the turn-on, automatic-refresh toggle button was pressed.)

Tip: Alternatively, let the toggle setting persist, and refresh as needed by pressing the F5 key.

Keyboard navigation

Most of the features of the IBM Spectrum Control GUI are accessible by using the keyboard. For those features that are not accessible, equivalent function is available by using the command-line interface (CLI), except as noted in the product release notes.

You can use keys or key combinations to perform operations and initiate many menu actions that can also be done through mouse actions. The following sections describe the keys or key combinations for different parts of the GUI:

For navigating in the GUI and the context-sensitive help system:

- To navigate to the next link, button, or topic within a panel, press Tab.
- To move to the previous link, button, or topic within a panel, press Shift+Tab.
- To select an object, when the object is in focus, press Enter.

For actions menus:

• To navigate to the grid header, press Tab.

- To reach the drop-down field, press the Left Arrow or Right Arrow key.
- To open the drop-down menu, press Enter.
- To select the menu items, press the Up Arrow or Down Arrow key.
- To start the action, press Enter.

For filters:

To specify a filter option and text:

- 1. Press Tab to navigate to the magnifying glass icon.
- 2. Press the Up Arrow or Down Arrow key to navigate the filtering list.
- 3. Press Enter to select a filtering option.
- 4. When a filtering option is selected, the cursor moves to the filter text box. Type the filter text and press Enter. To reset a filter, press Enter.

For text fields:

- To navigate to text fields, press Tab.
- To navigate to the fields that are available for editing, press Tab.
- To navigate to the next field or to the **Submit** button, press Tab.

For tables or lists:

- To navigate between column headers, focus on a column header and use the Left Arrow and Right Arrow keys to move to other column headers.
- To navigate between data cells, focus on a data cell and use the Left, Right, Up, Down, Pageup, and Pagedown Arrow keys.
- To sort a column, focus on a column header and press Enter. The focus remains on the column header after the sort occurs.
- To change the size of a column, focus on the column header, hold Shift+Control, and press the Left or Right Arrow keys.
- To follow a link in a data cell, focus on a data cell and press Shift+F9.
- To open a menu for a table row, focus on the row and press Shift+F10.
- To select consecutive rows, select the first row and hold Shift, press the Up or Down Arrow keys to go to the last row in the range, and press the Space bar to add the new rows to the selection.
- To select non-consecutive rows, select a row and hold Control, press the Up or Down Arrow keys, and press the Space bar to add the new row to the selection.

Restriction: For Chinese languages, the keyboard combination Control+Space bar is not enabled for selecting multiple rows at the same time.

Keyboard navigation with Firefox for Mac users: If you're using Firefox on a Mac with IBM Spectrum Control and want to use keyboard navigation, complete the following steps:

- 1. In Firefox, go to **Preferences** > **Advanced** > **General** and clear the check mark for **Always use the cursor keys to navigate within pages**. This step enables the use of Tab key to navigate between GUI elements.
- 2. In the URL address bar of Firefox, type about:config and press Enter.

Tip: If a warning prompt is displayed, click the button to accept the risk of changing browser settings. Existing settings won't be changed; instead, you'll be adding a preference setting for accessibility.

- 3. To add an accessibility preference for tab focus, right-click on the configuration page and select **New** > **Integer**.
- 4. In the New integer value window, type accessibility.tabfocus and click OK.
- 5. Type 7 to set the integer value and click **OK**.
- 6. Open your Mac's System Preferences app, go to **Keyboard > Shortcuts**, and select **All Controls**.

IBM and accessibility

For more information about IBM's commitment to accessibility, see the IBM Human Ability and Accessibility Center website at http://www.ibm.com/able.

Legal notices

This information was developed for products and services offered in the U.S.A. This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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. 19-21, Nihonbashi-Hakozakicho, Chuo-ku Tokyo 103-8510, 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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

COPYRIGHT LICENSE: This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details in the section entitled "Cookies, Web Beacons and Other Technologies," and the "Privacy on the IBM Cloud" at http://www.ibm.com/software/info/product-privacy.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy and http://www.ibm.com/pri

www.ibm.com/privacy/details/us/en/ in the section entitled "Cookies, Web Beacons and Other Technologies".

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corporation, 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 https://www.ibm.com/legal/us/en/copytrade.shtml.

Intel, Intel logo, Intel Xeon, and Pentium are trademarks or registered trademarks 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.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Red Hat® is a registered trademark of Red Hat, Inc. or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

VMware, the VMware logo, VMware Cloud Foundation, VMware Cloud Foundation Service, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Glossary

This glossary provides terms and definitions for IBM Spectrum Control.

The following cross-references are used in this glossary:

- See refers you from a nonpreferred term to the preferred term or from an abbreviation to the spelledout form.
- See also refers you to a related or contrasting term.

"A" on page 521 "B" on page 522 "C" on page 522 "D" on page 524 "E" on page 525 "F" on page 526 "G" on page 528 "H" on page 528 "I" on page 529 "J" on page 530 "K" on page 530 "L" on page 530 "M" on page 530 "N" on page 531 "O" on page 531 "P" on page 532 "R" on page 533 "S" on page 534 "T" on page 537 "U" on page 537 "V" on page 538 "W" on page 539 "Z" on page 539

A

activation key

See license key.

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.

advisory lock

A type of lock that a process holds on a region of a file that signals any other process to not use or lock the region or an overlapping region. Other processes are not forced to comply.

allocatable extent limit

A maximum total capacity for the system. The allocatable extent limit is calculated from pool extent sizes.

application key

See license key.

array

An ordered collection, or group, of physical devices (disk drive modules) that are used to define logical volumes or devices. An array is a group of drives designated to be managed with a Redundant Array of Independent Disks (RAID).

asynchronous replication

A type of replication in which control is given back to the application as soon as the write operation is made to the source volume. Some time later, the write operation is made to the target volume. See also synchronous replication.

audit log

An unalterable record of all commands or user interactions that are issued to the system.

authenticated user

A user who has logged in to the system with a valid account (user ID and password).

authentication

The mechanism by which a system determines what permissions a particular authenticated user has to access specific resources or actions. See also <u>authorization</u>.

authorization

The mechanism by which a system determines what permissions a particular authenticated user has to access specific resources or actions. See also authentication.

authorization code

An alphanumeric code generated for administrative functions, such as password resets or two-factor authentication bypass.

available capacity

The amount of usable capacity that is not yet used in a system, pool, array, or MDisk.

B

block storage

A unit of data storage on a device.

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.

C

cache

Storage or memory that is used to improve access times to instructions, data, or both. For example, data that resides in cache memory is normally a copy of data that resides elsewhere in slower, less expensive storage, such as on a disk or on another network node.

cache eviction

A process by which data associated with a file is removed from the cache system. The data is removed either by using a Least Recently Used (LRU) algorithm when configured General Parallel File System (GPFS) hard or soft quota limits are exceeded or by issuing a command. When referenced again in the cache system, the data that is associated with the file is retrieved from the home system.

caching I/O group

The I/O group in the system that performs the cache function for a volume.

call home

A communication link established between a product and a service provider. The product can use this link to place a call to a service provider when it requires service. With access to the machine, service personnel can perform service tasks, such as viewing error and problem logs or initiating trace and dump retrievals.

capacity

The amount of data that can be contained on a storage medium.

capacity recycling

The amount of provisioned capacity that can be recovered without causing stress or performance degradation. This capacity identifies the amount of resources that can be reclaimed and provisioned to other objects in an environment.

capacity threshold

The percent of total usable physical capacity that used capacity must exceed before a notification is sent. See also total usable physical capacity.

certificate

A digital document that binds a public key to the identity of the certificate owner, thereby enabling the certificate owner to be authenticated. A certificate is issued by a certificate authority and is digitally signed by that authority.

change volume

A volume that is used in Global Mirror that holds earlier consistent revisions of data when changes are made.

child pool

A user-defined capacity that is formed from capacity that is defined either in another pool or a system. See also parent pool.

CIFS

See Common Internet File System.

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).

CIM object manager (CIMOM)

The common conceptual framework for data management that receives, validates, and authenticates the CIM requests from the client application. It then directs the requests to the appropriate component or service provider.

CIMOM

See CIM object manager.

CKD

See count key data.

CKD record

See count-key-data record.

CLI

See command-line interface.

client

A software program or computer that requests services from a server. See also host, server.

cloud account

An agreement with a cloud service provider to use storage or other services at that service provider. Access to the cloud account is granted by presenting valid credentials.

cluster

- 1. A collection of complete systems that work together to provide a single, unified computing capability.
- 2. A group of computers and other resources that operate together as a single system.
- 3. A loosely coupled collection of independent systems (or nodes) organized into a network for the purpose of sharing resources and communicating with each other.
- 4. 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.
- 5. In Storwize® V7000, a pair of nodes that provides a single configuration and service interface.

command-line interface (CLI)

A computer interface in which the input and output are text based.

Common Information Model (CIM)

An implementation-neutral, object-oriented schema for describing network management or systems management information. The Distributed Management Task Force (DMTF) develops and maintains CIM specifications.

Common Internet File System (CIFS)

A protocol that manages shared, remote file access for applications to files, printers, serial ports, and so on over a TCP/IP network.

community name

The part of an SNMP message that represents a password-like name and that is used to authenticate the SNMP message.

compression

A function that removes repetitive characters, spaces, strings of characters, or binary data from the data being processed and replaces characters with control characters. Compression reduces the amount of storage space that is required for data.

compute node

An independent machine that contains one or more microprocessors, memory, storage, and network controllers and runs its own operating system and applications.

concurrent copy

A function of the DFSMSdss component that is used to back up any collection of data at a point in time with minimum down time for the database or application that uses the collection of data.

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.

copyback

A process that moves data back to its expected or preferred location to maintain an array in a more efficient configuration after a failed drive is replaced.

copy set

The set of source volumes or target volumes involved in a FlashCopy operation.

count key data

- 1. An architecture for a direct access storage device (DASD) device or logical device that specifies the access mechanisms for the logical data units on the device through a specific set of supported channel commands. Extensions to the CKD command set form the basis of Extended CKD.
- 2. A data recording format that uses self-defining record formats in which each record on a volume 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. See also data record, storage architecture type.

count-key-data record (CKD record)

See data record.

CRU

See customer-replaceable unit.

customer-replaceable unit (CRU)

An assembly or part that can be replaced in its entirety by a user when any one of its components fails.

cylinder

A unit of storage on a count-key-data (CKD) device with a fixed number of tracks.

D

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 <u>discovery</u>.

data consistency

A characteristic of the data at the target site where dependent write order is maintained to guarantee the recoverability of applications.

data record

A basic unit of data recording format. See also count key data, fixed-block architecture.

data reduction

A set of techniques that can be used to reduce the amount of usable capacity that is required to store data. Examples of data reduction include data deduplication and compression. See also <u>data</u> reduction savings, stored capacity.

data reduction savings

The total amount of usable capacity that is saved in a system, pool, or volume through the application of an algorithm such as compression or deduplication on the written data. This saved capacity is the difference between the written capacity and the used capacity. See also data reduction.

data source

A storage resource or agent that provides data about a storage environment.

destage

To move data from cache to a nonvolatile storage medium.

discovery

The process of finding resources within an enterprise, including finding the new location of monitored resources that were moved. See also data collection.

distributed RAID

An alternative RAID scheme where the number of drives that are used to store the array can be greater than the equivalent, typical RAID scheme. The same data stripes are distributed across a greater number of drives, which increases the opportunity for parallel I/O and hence improves overall array performance. See also rebuild area.

DNS

See Domain Name System.

Domain Name System (DNS)

The distributed database system that maps domain names to IP addresses.

drive

A data storage device. A drive can be either a magnetic disk drive or a solid-state drive (SSD).

drive class

A combination of drive technology and speed, which uniquely defines a class of drives that have approximately the same performance characteristics.

drive technology

A category of a drive that pertains to the method and reliability of the data storage techniques being used on the drive. Possible values include enterprise (ENT) drive, nearline (NL) drive, or solid-state drive (SSD).

Ε

ECKD

See extended count key data.

effective capacity

The amount of provisioned capacity that can be created in a system or pool without running out of usable capacity given the current data reduction savings being achieved. This capacity equals the usable capacity divided by the data reduction savings percentage.

enclosure

The metal structure in which various electronic components are mounted.

encryption deadlock

The inability to access encryption keys to decrypt data. See also encryption recovery key.

encryption key label

The list of encryption key labels used by the storage system to identify keys that will be used on the key server.

encryption key manager

See encryption key server.

encryption key server

An internal or external system that runs a key manager that receives and then serves existing encryption keys or certificates to a storage system.

encryption recovery key

An encryption key that allows a method to recover from an encryption deadlock situation where the normal encryption key servers are not available. See also encryption deadlock.

enterprise

Pertaining to a type of data storage device that has higher error recovery limits, vibration tolerance, and end-to-end error detection than standard desktop hard drives.

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 the data server using Java Database Connectivity (JDBC).

event

An occurrence of significance to a task or system. Events can include completion or failure of an operation, a user action, or the change in state of a process.

extended count key data (ECKD)

An extension of the count-key-data (CKD) architecture. It includes additional commands that can be used to improve performance.

extent type

See storage architecture type.

F

failback

The restoration of an appliance to its initial configuration after detection and repair of a failed network or component.

failover

An automatic operation that switches to a redundant or standby system or node in the event of a software, hardware, or network interruption.

FB

See fixed block.

FBA

See fixed-block architecture.

FC

See Fibre Channel.

FC-AL

See Fibre Channel Arbitrated Loop.

FCIP

See Fibre Channel over IP.

FCP

See Fibre Channel Protocol.

feature activation code

See license key.

Fibre Channel (FC)

A technology for transmitting data between computer devices. It is especially suited for attaching computer servers to shared storage devices and for interconnecting storage controllers and drives. See also zoning.

Fibre Channel Arbitrated Loop (FC-AL)

An implementation of the Fibre Channel standards that uses a ring topology for the communication fabric; refer to American National Standards Institute (ANSI) INCITS 272-1996, (R2001). In this topology, two or more Fibre Channel end points are interconnected through a looped interface.

Fibre Channel connection (FICON)

A Fibre Channel communication protocol designed for IBM mainframe computers and peripherals.

Fibre Channel extender

A device used to extend a Fibre Channel link over a greater distance than is supported by the standard, usually a number of miles or kilometers. Devices must be deployed in pairs at each end of a link.

Fibre Channel over IP (FCIP)

A network storage technology that combines the features of the Fibre Channel Protocol and the Internet Protocol (IP) to connect distributed SANs over large distances.

Fibre Channel Protocol (FCP)

The serial SCSI command protocol used on Fibre Channel networks. See also open system.

FICON

See Fibre Channel connection.

field-replaceable unit (FRU)

An assembly that is replaced in its entirety when any one of its components fails.

file module

A component that provides file systems to network users. A file module must be provided with storage for the file systems.

fileset

See file set.

file set

- 1. A subset of a file system that provides granularity of functions such as snapshots or quotas within the file system.
- 2. A hierarchical grouping of files managed as a unit for balancing workload across a cluster.

file system (FS)

A collection of files and certain attributes associated with those files.

file system storage

Data storage that is organized into files and directories.

fixed block (FB)

See fixed-block architecture.

fixed-block architecture (FBA)

An architecture for a virtual device that specifies the format of and access mechanisms for the virtual data units on the device. The virtual data unit is a block. All blocks on the device are the same size (fixed size). The system can access them independently. See also <u>data record</u>, <u>storage architecture</u> type.

FlashCopy

- 1. Pertaining to a point-in-time copy where a virtual copy of a volume is created. The target volume maintains the contents of the volume at the point in time when the copy was established. Any subsequent write operations to the source volume are not reflected on the target volume.
- 2. An optional feature of the Storage System DS family that can make an instant copy of data, that is, a point-in-time copy of a volume.

flash drive

A data storage device, which is typically removable and rewritable, that uses solid-state memory to store persistent data. See also flash module.

flash module

A modular hardware unit containing flash memory, one or more flash controllers, and associated electronics. See also flash drive.

flush-through mode

See write-through mode.

form factor

The industry-standard physical dimensions of a storage system drive enclosure. Possible values include "3.5 inch", "2.5 inch", and "1.8 inch."

frame

The hardware support structure, covers, and all electrical parts mounted therein that are packaged as one entity for shipping.

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.

FRU

See field-replaceable unit.

FS

See file system.

full restore operation

A copy operation where a local volume is created by reading an entire a volume snapshot from cloud storage.

full snapshot

A type of volume snapshot that contains all the volume data. When a full snapshot is created, an entire copy of the volume data is transmitted to the cloud.

G

General Parallel File System (GPFS)

A high-performance shared-disk file system that can provide data access from nodes in a clustered system environment.

Global Copy

A non-synchronous long-distance copy option for data migration and backup. See also <u>remote mirror</u> and copy.

globally unique identifier (GUID)

An algorithmically determined number that uniquely identifies an entity within a system.

Global Mirror

A method of an asynchronous replication that maintains data consistency across multiple volumes within or across multiple systems. Global Mirror is generally used where distances between the source site and target site cause increased latency beyond what the application can accept.

GPFS

See General Parallel File 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.

Н

Hardware Management Console (HMC)

- 1. In a system storage environment, a system that acts as the focal point for configuration, management of Copy Services functions, and maintenance.
- 2. A system that controls managed systems, including the management of logical partitions and use of Capacity Upgrade on Demand. Using service applications, the HMC communicates with managed systems to detect and consolidate information, which can then be sent for analysis.

heat map

An overlay that shows where the user performed actions on a web page. The data is summarized and shown as a colored marker at the location of the user action over a snapshot of the web page. The overlay lets the user see where and how the web page user interacted with the web page.

HMC

See Hardware Management Console.

host

A physical or virtual computer system that hosts computer applications, with the host and the applications using storage. See also client, host, server.

host cluster

A configured set of physical or virtual hosts that share one or more storage volumes in order to increase scalability or availability of computer applications.

host interface card

See interface card.

host object

A logical representation of a host within a storage system that is used to represent the host for configuration tasks.

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.

hot-spare

Pertaining to redundant hardware (such as an adapter, a disk, a drive, or a server) that is installed and available in the event of a hardware failure.

HyperSwap

Pertaining to a function that provides continuous, transparent availability against storage errors and site failures, and is based on synchronous replication.

hypervisor

Software or a physical device that enables multiple instances of operating systems to run simultaneously on the same hardware.

Ι

incremental restore operation

A copy operation where a local volume is modified to match a volume snapshot by reading from cloud storage only the parts of the volume snapshot that differ from the local volume.

incremental snapshot

A type of volume snapshot where the changes to a local volume relative to the volume's previous snapshot are stored on cloud storage.

input/output (I/O)

Pertaining to a device, process, channel, or communication path involved in data input, data output, or both.

interface card

An optional part of a node canister that provides the system with additional host and storage connectivity options.

interface node

A node that connects a system to an Internet Protocol (IP) network for file-serving capabilities by using service protocols.

Internet Small Computer System Interface (iSCSI)

An IP-based standard for linking data storage devices over a network and transferring data by carrying SCSI commands over IP networks. See also Small Computer System Interface.

I/O

See input/output.

I/O enclosure

A hardware unit in a storage system where data is transferred into and out of the system.

iSCSI

See Internet Small Computer System Interface.

J

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.

K

key server

- 1. A server that negotiates the values that determine the characteristics of a dynamic virtual private network (VPN) connection that is established between two endpoints.
- 2. See encryption key server.

L

launch-in-context

An operation in which a user starts a secondary application from a primary application to perform a specific task. Using the parameters, navigation instructions, and user credentials that are supplied by the primary application, the secondary application opens to the specific place in which to complete the task.

licensed capacity

The amount of capacity on a storage system that a user is entitled to configure.

license key

An alphanumeric code that activates a licensed function on a product.

license key file

A file that contains one or more licensed keys.

logical unit number (LUN)

In the Small Computer System Interface (SCSI) standard, a unique identifier used to differentiate devices, each of which is a logical unit (LU).

LUN

See logical unit number.

M

machine signature

A string of characters that identifies a system. A machine signature might be required to obtain a license key.

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 node

A node that is used for configuring, administering, and monitoring a 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.

maximum replication delay

The number of seconds that Metro Mirror or Global Mirror replication can delay a write operation to a volume.

Metro Global Mirror

- A cascaded solution where Metro Mirror synchronously copies data to the target site. This Metro
 Mirror target is the source volume for Global Mirror that asynchronously copies data to a third site.
 This solution has the potential to provide a disaster recovery with no data loss at Global Mirror
 distances when the intermediate site does not participate in the disaster that occurs at the
 production site.
- 2. 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.

Metro Mirror

A method of synchronous replication that maintains data consistency across multiple volumes within the system. Metro Mirror is generally used when the write latency caused by the distance between the source site and target site is acceptable to application performance.

MIB

See Management Information Base.

Ν

namespace

The scope within which a Common Information Model (CIM) schema applies.

native interface

An interface that is specific to a system or subsystem.

nearline

Pertaining to a type of storage in which data is available in a short amount of time, but not instantly.

nearline SAS drive

A drive that combines the high capacity data storage technology of a Serial Advanced Technology Attachment (SATA) drive with the benefits of a serial-attached SCSI (SAS) interface for improved connectivity.

node

A single processing unit within a system. For redundancy, multiple nodes are typically deployed to make up a system.

0

open system

A system that complies with industry-defined interoperability standards. An open system can be connected to other systems complying with the same standards. See also <u>Fibre Channel Protocol</u>, Small Computer System Interface.

order confirmation code

See authorization code.

overhead capacity

An amount of usable capacity that is occupied by metadata in a system or pool and other data that is used for system operations.

overprovisioned ratio

The ratio of provisioned capacity to usable capacity in a system or pool.

overprovisioning

The result of creating more provisioned capacity in a storage system or pool than there is usable capacity. Overprovisioning occurs when thin provisioning or data reduction techniques ensure that the used capacity of the provisioned volumes is less than their provisioned capacity.

P

parent pool

A storage pool that receives its capacity from MDisks and has, or will have, some of its capacity allocated to child pools. See also child pool.

performance group

A collection of volumes that is assigned the same performance characteristics. See also <u>performance</u> policy.

performance policy

A policy that specifies performance characteristics, for example quality of service (QoS). See also performance group.

PFC

See priority flow control.

ping

- 1. The command that sends an Internet Control Message Protocol (ICMP) echo-request packet to a gateway, router, or host with the expectation of receiving a reply.
- 2. 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.

pool

- 1. See storage pool.
- 2. A grouping of storage space that consists of volumes, logical unit numbers (LUNs), or addresses that share a common set of administrative characteristics.

pool pair

Two storage pools that are required to balance workload. Each storage pool is controlled by a separate node.

port

The physical entity within a host, system, or storage system that performs the data communication (transmitting and receiving) over the Fibre Channel.

practice volume

A volume that can be used to test disaster-recovery actions while maintaining disaster-recovery capability.

primary site

A physical or virtual site that is made up of hardware, network, and storage resources. Typically, production operations run at the primary site. Data can be replicated to a secondary site for disaster recovery and failover operations. See also secondary site.

primordial pool

Storage capacity that is unallocated on a storage device. Storage pools are created by allocating storage capacity from primordial pools.

priority flow control (PFC)

A link-level flow control mechanism, IEEE standard 802.1Qbb. PFC operates on individual priorities. Instead of pausing all traffic on a link, PFC is used to selectively pause traffic according to its class.

probe

A data collection job that itemizes and creates an inventory of assets, such as computers, controllers, disk drives, file systems, and logical units.

projected capacity

The estimated volume capacity that is available for volume creation, given the current average performance of any data compression, excluding thin-provisioning savings. See also thin-provisioning savings.

protocol

A set of rules controlling the communication and transfer of data between two or more devices or systems in a communication network.

provisioned capacity

The total capacity of all volumes and volume copies in a system or pool.

R

rack

A free-standing structure that can hold multiple servers, storage systems, chassis, switches, and other devices.

RAID

See Redundant Array of Independent Disks.

RAID 0

A data striping technique, which is commonly called RAID Level 0 or RAID 0 because of its similarity to common, RAID, data-mapping techniques. It includes no data protection, however, so, strictly speaking, the appellation RAID is a misnomer. RAID 0 is also known as data striping.

PATD 1

A form of storage array in which two or more identical copies of data are maintained on separate media.

RAID 10

A collection of two or more physical drives that present to the host an image of one or more drives. In the event of a physical device failure, the data can be read or regenerated from the other drives in the RAID due to data redundancy.

RAID 5

A form of parity RAID in which the disks operate independently, the data stripe size is no smaller than the exported block size, and parity check data is distributed across the array's disks.

RAID 6

A form of RAID that can continue to process read and write requests to all of an array's virtual disks in the presence of two concurrent disk failures.

RAID level

The level of protection provided by the specific techniques of striping, mirroring, or parity used by a Redundant Array of Independent Disks (RAID).

RAID type

See RAID level.

raw capacity

The reported capacity of the drives in the system before formatting or RAID (Redundant Array of Independent Disks) is applied.

rebuild area

Reserved capacity that is distributed across all drives in a redundant array of drives. If a drive in the array fails, the lost array data is systematically restored into the reserved capacity, returning

redundancy to the array. The duration of the restoration process is minimized because all drive members simultaneously participate in restoring the data. See also distributed RAID.

reclaimable capacity

The amount of provisioned capacity that can be recovered without causing stress or performance degradation. This capacity identifies the amount of resources that can be reclaimed and provisioned to other objects in an environment.

reclaimed capacity

See reclaimable capacity.

recovery key

See encryption recovery key.

recovery point objective

The maximum amount of data loss that can be tolerated during a service interruption.

Redundant Array of Independent Disks (RAID)

A collection of two or more physical disk drives that present to the host an image of one or more logical disk drives. In the event of a physical device failure, the data can be read or regenerated from the other disk drives in the array due to data redundancy.

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 Global Copy.

repo

See repository.

repository (repo)

A persistent storage area for data and other application resources.

reserved capacity

The amount of used capacity that is made up of capacity reserved for system use. See also <u>total</u> usable physical capacity.

resource

In a storage environment, an entity that is monitored. Resources can include fabrics, switches, computers, and storage systems.

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).

S

scan

A data collection job that monitors storage usage and file statistics on the resources in an environment.

SCSI

See Small Computer System Interface.

SCSI device

A product, such as a drive or adapter, connected to a host through an I/O interface using the Small Computer System Interface (SCSI) protocol. A SCSI device is either an initiator, target, or both. See also Small Computer System Interface.

SCSI-FCP

See SCSI Fibre Channel Protocol.

SCSI Fibre Channel Protocol (SCSI-FCP)

A standard that defines the protocol used to transfer Small Computer System Interface (SCSI) commands over the transport physical layer of the Fibre-Channel interface. This standard is published by ANSI as X3.269-1996.

SCSI initiator

The system component that initiates communications with attached targets.

SCSI target

A device that acts as a subordinate to a SCSI initiator and consists of a set of one or more logical units (LUs), each with an assigned logical unit number (LUN). The LUs on the SCSI target are typically I/O devices.

secondary site

A physical or virtual site that is made up of the hardware, network, and storage resources that support the recovery needs of the primary site. When a failure occurs at the primary site, operations can continue at the secondary site. See also primary site.

server

A computer program or a device that provides functions for other programs or devices, called clients. See also client, host.

Server Message Block (SMB)

A protocol that manages requests and responses in a client/server environment so that clients on a network can share files, directories, and devices. See also Server Message Block 2.0.

Server Message Block 2.0

A higher performing, more scalable version of Server Message Block (SMB). This protocol can send multiple commands in the same packet and uses larger buffer sizes. See also Server Message Block.

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.

Small Computer System Interface (SCSI)

An ANSI-standard electronic interface that allows personal computers to communicate with peripheral hardware, such as disk drives, tape drives, CD-ROM drives, printers, and scanners faster and more flexibly than previous interfaces. See also <u>Internet Small Computer System Interface</u>, <u>open system</u>, SCSI device.

SMB

See Server Message Block.

SMI-S

See Storage Management Initiative Specification.

solid-state drive (SSD)

- 1. A storage device that contains nonvolatile flash memory. A solid-state drive (SSD) has no moving mechanical components.
- 2. See flash drive.

space

See capacity.

space efficient

See thin provisioning.

spare drive

A drive reserved in an array for rebuilding a failed drive in a RAID. Should a drive fail in a RAID, a spare drive from within that device adapter (DA) pair will be selected to rebuild it.

SSD

See solid-state drive.

standard-provisioned volume

A volume that completely uses storage at creation.

standard provisioning

The ability to completely use a volume's capacity for that specific volume.

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 architecture type (storage type)

The type of storage architecture, either count key data (CKD) or fixed block (FB), for which an array, pool, or volume is provisioned. See also count key data, fixed-block architecture.

storage enclosure

A specialized chassis that is designed to hold and power drives while providing a mechanism to allow them to communicate to one or more separate computers.

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 node

A component of a storage system that provides internal storage or a connection to one or more external storage systems.

storage pod

- 1. A subcomponent of a network-attached storage (NAS) system that consists of two or more storage nodes and one or more supported storage systems.
- 2. A logical entity of components of a system consisting of two storage nodes, and one or two storage subsystems directly connected with these storage nodes.

storage pool (pool)

A collection of storage that identifies an underlying set of resources. These resources provide the capacity and management requirements for a volume or set of volumes.

storage resource group

A named collection of logically related resources that are monitored by IBM Spectrum Control. Monitored resources can include fabrics, switches, computers, storage systems, and other storage resource groups.

storage system

A system that provides persistent storage within a network. A storage system can include facilities for host attachment, user role authentication, a command-line interface (CLI), a graphical user interface (GUI), and storage devices that most often include Redundant Array of Independent Disks (RAID) controllers. It might also include agents for enabling third-party management software to monitor or manage the storage devices.

storage type

See storage architecture type.

stored capacity

The amount of capacity that is used to store data that is written by a host after data reduction See also data reduction, total usable physical capacity.

support assistance

A function that is used to provide support personnel access to the system to complete troubleshooting and maintenance tasks.

synchronous replication

A type of replication in which the application write operation is made to both the source volume and target volume before control is given back to the application. See also <u>asynchronous replication</u>.

syslog

A standard for transmitting and storing log messages from many sources to a centralized location to enhance system management.

T

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.

thin-provisioned volume

A volume that allocates storage when data is written to it.

thin provisioning

- 1. The ability to defer capacity allocation on a storage resource until data is actually written to it.
- 2. A mechanism that provides the ability to define logical volume sizes that are larger than the physical capacity installed on the system.

thin-provisioning savings

The total amount of usable capacity that is saved in a system, pool, or volume by consuming usable capacity only when needed as a result of write operations. The capacity that is saved is the difference between the provisioned capacity minus the written capacity. See also projected capacity, volume capacity, written capacity.

total capacity savings

The total amount of usable capacity that is saved in a system, pool, or volume through thin provisioning and data reduction techniques. This saved capacity is the difference between the used usable capacity and the provisioned capacity.

total usable physical capacity

The amount of physical configured storage space that is available for stored capacity or reserved capacity. This capacity can consist of both internal storage through arrays and external storage through MDisks. See also capacity threshold, reserved capacity, stored capacity.

track space-efficient volume (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.

transparent cloud tiering

The functions that use cloud storage as an extension of on-premises storage.

trial license

A temporary entitlement to use a licensed function.

TSE for FlashCopy

A thin-provisioning method in which storage space is allocated from a TSE repository on an as needed basis. See also TSE repository.

TSE repository

The amount of capacity in a storage pool reserved for volumes that use a thin-provisioning method of TSE for FlashCopy. See also TSE for FlashCopy.

TSE volume

See track space-efficient volume.

U

unmapped volume capacity

The amount of volume capacity that is not mapped to a host. See also volume capacity.

update

- 1. To apply fixes to a system.
- 2. Software maintenance such as a manufacturing refresh, refresh pack, or fix pack that changes the modification level of a product.
- 3. To modify a file or data set with current information.

upgrade

- 1. Any hardware or software change to a later release, or any hardware addition or software addition.
- 2. To install a new version or release of a product to replace an earlier version or release of the same product.

usable capacity

The amount of capacity that is provided for storing data on a system, pool, array, or MDisk after formatting and RAID techniques are applied.

used capacity

The amount of usable capacity that is taken up by data or overhead capacity in a system, pool, array, or MDisk after data reduction techniques have been applied.

user role

An identifier that is assigned to a user that defines the set of permissions that are granted to that user.

V

virtual capacity

See provisioned capacity.

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.

virtualized capacity

The amount of capacity that is contributed to a storage pool by a given provisioning group.

virtual machine (VM)

An emulation of a particular computer system. Virtual machines operate based on the computer architecture and functions of a real or hypothetical computer. Their implementations might involve specialized hardware, software, or a combination of both.

virtual storage area network (VSAN)

A fabric within the storage area network (SAN).

VM

See virtual machine.

volume

A fixed amount of physical or virtual storage on a data storage medium.

volume access set

The set of I/O groups that allows host access to a volume. This set can optionally include the caching I/O group.

volume capacity

The total capacity for all volumes in a system or storage pool. Volume capacity is defined by the client when a volume is created and surfaced to the host. See also <u>thin-provisioning savings</u>, <u>unmapped</u> volume capacity.

volume snapshot

A collection of objects on a cloud storage account that represents the data of a volume at a particular time.

VSAN

See virtual storage area network.

W

worldwide ID (WWID)

A name identifier that is unique worldwide and that is represented by a 64-bit value that includes the IEEE-assigned organizationally unique identifier (OUI).

worldwide name (WWN)

A 64-bit, unsigned name identifier that is unique.

worldwide node name (WWNN)

A unique 64-bit identifier for a host containing a Fibre Channel port. See also worldwide port name.

worldwide port name (WWPN)

A unique 64-bit identifier associated with a Fibre Channel adapter port. The WWPN is assigned in an implementation-independent and protocol-independent manner. See also worldwide node name.

write-through mode

A process in which data is written to a storage device at the same time as the data is cached.

written capacity

The amount of usable capacity that would have been used to store written data in a system or pool if data reduction was not applied. See also thin-provisioning savings.

written capacity limit

The largest amount of capacity that can be written to a drive, array, or MDisk. The limit can be reached even when usable capacity is still available.

WWID

See worldwide ID.

WWN

See worldwide name.

WWNN

See worldwide node name.

WWPN

See worldwide port name.

Z

z Global Mirror

A method of an asynchronous replication function that maintains data consistency across multiple volumes that are attached to a z/OS system. Time-based data consistency is maintained through the Data Facility Storage Management Subsystem (DFSMS) system data mover (SDM) component.

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 to be managed together.

zone set

A group of zones that function together on a fabric.

zoning

The grouping of multiple ports to form a virtual, private, storage network. Ports that are members of a zone can communicate with each other, but are isolated from ports in other zones. See also <u>Fibre</u> Channel.

Index

A	analysis period <i>(continued)</i>
	performance data, change <u>449</u>
about this document <u>xi</u>	analyze tiering
accessibility features	analysis period, change 449
for IBM Spectrum Control 513	conditions, setting 88
active data 364	hypervisors 442
activity levels	servers 442
formula 446	storage pools 443
activity limits	storage virtualizers 442
formula 446	thresholds, modify 444
adding	volume placement 88
NetApp resources 124	volumes 443
adding applications 409	volumes, collocate <u>90</u>
adding hypervisors	announcements <u>62</u>
ESX and ESXi <u>132</u>	application
vCenter Server system <u>132</u>	scenario <u>94</u>
adding resources	tutorial <u>94</u>
fabrics <u>129</u>	application context <u>405</u>
hypervisors <u>132</u> , <u>133</u>	applications
rollup servers <u>136</u>	adding applications to departments for monitoring 409
servers with agents 134	application hierarchy 399
storage systems	creating 399
Hitachi 121	creating applications for monitoring 400
switches 129	custom capacity reports 487, 488
vCenter Server system 133	custom report, create 487
adding subdepartments 408	defining alerts for attribute and capacity changes 174
administering general groups 417	defining alerts for performance metrics 176
agents	monitoring 398
CIM 84	removing 405
alert conditions and violations 75, 151, 152	removing applications for monitoring 405
alert notification settings	resources 405
defining 157	scenarios
alert notifications	
email 186	viewing capacity usage trends <u>96</u>
	tutorials
Netcool/OMNIbus <u>187</u>	viewing capacity usage trends 96
SNMP 187	viewing information 398
alert policies	viewing status 398
viewing and administering <u>156</u>	applications and departments
alerting	tutorials
alert log <u>75</u> , <u>151</u> , <u>152</u>	monitoring capacity usage 91
defining alerts for attribute and capacity changes <u>159</u> ,	architecture
174	description <u>42</u>
defining alerts for performance metrics 176	audience <u>ix</u>
defining general group alerts for attribute and capacity	
changes <u>180</u>	В
notifications <u>75</u> , <u>151</u> , <u>152</u>	
alerts	balance pools
actions <u>226</u>	criteria 446
configuring for performance monitors 147	modify 445
configuring for probes 142	volumes, collocate 90
notifications 226	Balance Pools 445
triggering conditions 188	Beta program
viewing and administering 154, 155	
viewing violations of capacity threshold 367	advantages <u>xxxii</u>
viewing violations of performance thresholds 275	overview <u>xxxii</u>
alias 82	block storage
analysis period	capacity matrice 367
	capacity metrics 367

blogs 17, 19	capacity view
	accessing <u>355</u>
C	applications <u>358</u>
	bookmarks <u>364</u>
capacity	containers <u>358</u>
adding to a capacity pool 426	controls <u>356</u>
alerts 367	file system pools <u>358</u>
block storage systems 348	file systems 358
capacity view 355	filesets 358
containers 348	metrics 356, 358
controls for the capacity view 356	pools 358
file storage systems 348	storage systems 358
- ·	tiers 358
file system pools 348	volumes 358
file systems 348	capacity-to-limit 350
filesets 348	capacity)
metrics <u>356, 367</u>	used capacity growth rate 240
monitoring 348	chargeback reports 492, 494
non-tiered storage <u>440</u>	CIM agents 84
object storage systems <u>348</u>	Cisco MDS 9000 83
pool <u>240</u>	
pools <u>348</u>	cloud configuration 68, 421
removing from a capacity pool <u>427</u>	cloud object storage
tiered storage <u>440</u>	, vaults, accessibility risk <u>257</u>
tiers <u>348</u>	, vaults, inaccessible <u>256</u>
viewing threshold violations 367	cloud storage 364, 366
volumes 244, 348	collaborating <u>63</u>
capacity charts	collecting data
block pools 367	GPFS snapshots <u>150</u>
block storage systems 367	collection
block volumes 367	overview <u>64</u>
compliance check 350	comments, sending <u>xi</u>
containers 390	components <u>42</u>
object storage systems 390	computer alerts
tiers 391	triggering conditions <u>219</u>
capacity license model	condition of resources
category of drives 30, 34, 37	how the condition of resources is determined 248
uncategorized drives 30, 34, 37	condition of resources in the web-based GUI
capacity license model) 38	viewing the condition of top-level resources 252
capacity limit	conditions and violations 75, 151, 152
adjusted used capacity 351	configuring —— —— ——
capacity-to-limit 351	after an upgrade 37
	alerts for performance monitors 147
capacity limit, compliance 350	alerts for probes 142
capacity limit, remove <u>351</u>	license model 37
capacity metrics	license model after an upgrade 37
capacity charts 367	licensing 37
containers 390	NetApp resources 124
object storage systems 390	consumer reports
pools <u>367</u>	capacity 496
storage systems <u>367</u>	containers
tiers <u>391</u>	capacity charts 390
volumes <u>367</u>	capacity metrics 390
capacity planning 479	
capacity pool	controls for capacity view
adding resources to <u>426</u>	export capacity information 356
creating <u>426</u>	hide and show resources 356
deleting <u>428</u>	hide and show the legend 356
modifying 427	metrics 356
removing resources from 427	remove resources <u>356</u>
viewing capacity of 427	select metrics <u>356</u>
viewing details of 427	specify a time range <u>356</u>
viewing resources of 427	specify granularity <u>356</u>
capacity reports	view chart 356
applications 487, 488	view chart in a new window 356
capacity summary reports 497, 498	view table 356

controls for performance view	Dell EMC
export performance information <u>266</u>	disks, view performance metrics 321
hide and show resources 266	nodes, view performance metrics 321
hide and show the legend <u>266</u>	pools, view performance metrics 321
remove resources <u>266</u>	ports, view performance metrics 321
select metrics <u>266</u>	volumes, view performance metrics 321
specify a time range <u>266</u>	department
specify granularity <u>266</u>	tutorials
view chart <u>266</u>	comparing storage in each department 93
view chart in a new window 266	department context 411
view table <u>266</u>	departments
creating 400, 403, 407, 422	applications <u>411</u>
creating applications 400	creating departments for monitoring 407
creating departments <u>407</u>	monitoring <u>406</u>
creating subcomponents <u>403</u>	removing <u>411</u>
custom alert for resource	removing departments for monitoring 410
editing <u>164</u>	viewing information 406
custom alerts for applications	viewing status 406
defining 178	departments and applications
editing 178	scenarios
custom alerts for general groups	relational 96
defining 184	tutorials —
editing 184	different levels of business hierarchy 91
custom reports	details 404, 410
applications 487, 488	discover 150
capacity views, selecting time ranges 489	domain port zoning 82
delete 499	Domain_ID 82
edit 473, 499	
email, send 483, 484, 489	E
performance views, selecting time ranges 489	E
relative time ranges, selecting 489	EB, See EiB
save 484	EiB 79
schedule 483, 484, 489	EIF rules files 187
	email
D	capacity reports 487, 488
	email alert notifications 186
dashboard 233	email reports 475, 479, 482, 497
data collection	enclosure license model 30, 34, 37, 38
creating performance monitors in the web-based GUI	enclosures
144	viewing 244
modifying performance monitors in the web-based GUI	End of Support 62
146	enhancement request 63
performance monitors 144	ESX and ESXi
starting and stopping performance monitors in the GUI	adding 132
147	exported file systems 150
verifying that a performance monitor is running in the	exporting information in the GUI 452
web-based GUI <u>145</u>	exporting information in the web-based GUI 412
verifying that a probe is running in the web-based GUI	exportPerformanceData script 277
<u>140</u>	external storage <u>364</u> , <u>366</u>
viewing all performance monitors in theweb-based GUI	
148	F
viewing performance monitor logs in the web-based GUI	•
<u>149</u>	fabric 81
viewing performance monitors for a resource in theweb-	fabric alerts
based GUI 148	triggering conditions 218
viewing performance monitors in theweb-based GUI	fabrics
148	adding 129, 132
default	agent types for monitoring 132
zone <u>82</u>	exporting performance data 276
defining alerts for attribute and capacity changes 159, 174	modifying properties 450
defining application alerts for performance metrics 176	performance data, exporting 276
defining general group alerts for attribute and capacity	properties that can be modified 450
changes 180	viewing information in the GUI 231
defining general group alerts for performance metrics 182	

fabrics (continued) viewing status in the web-based GUI 247 FCID zoning 82 filtering lists in the web-based GUI 454 filters assign resources 401 sorting rows 455 fix packs 60	hierarchies (continued) viewing information 411 Hitachi VSP storage systems disks, view performance metrics 332 nodes, view performance metrics 332 pools, view performance metrics 332 ports, view performance metrics 332 volumes, view performance metrics 332 horizontal bar charts 238 host definition default 434
general group hierarchy 414 general groups add resources 414 defining alerts for attribute and capacity changes 180 defining alerts for performance metrics 182 subgroups, add 415 GiB 79 glossary 521 GPFS snapshots in IBM Spectrum Scale collecting data 150 groups general groups adding resources 416 administering actions 417 GUI	for fibre_channel ports 434 provisioning 434 how the condition of resources is determined 248 hypervisor alerts triggering conditions 210 hypervisors adding 132, 133 modifying properties 450 optimized volumes, collocate 90 properties that can be modified 450 tiering, analyze 442 user roles 116 viewing information in the GUI 231 viewing status in the web-based GUI 247
customizing lists 454 data collection performance monitors 144 starting performance monitors 147 stopping performance monitors 147 exporting information about applications and departments 412 exporting information about resources, tasks, and alerts 452 how information is organized 231 internal, related, and object resources, list 111 managing tasks 455 monitoring fabrics 82 monitoring switches 82 Opening the management GUI for a resource 452 resources that can be monitored 111 saving information about applications and departments 412 saving information about resources, tasks, and alerts 452 tasks managing provisioning tasks 459 managing tiering storage, balancing pools, and transforming-storage tasks 461 viewing all performance monitors 148 viewing information about internal, object, and related resources 237 viewing information about resources 231 viewing pending tasks 456 viewing tasks 456	I/O rate volumes, tiering 88 IBM Copy Services Manager 59 IBM Spectrum Control architecture 42 component overview 42 overview 41 IBM Spectrum Control REST API using 85, 468 IBM Spectrum Scale file systems, view performance metrics 317 nodes, view performance metrics 317 IBM Spectrum Virtualize host connections, view performance metrics 289 I/O groups, view performance metrics 289 key performance indicators 274 managed disks, view performance metrics 289 nodes, view performance metrics 289 pools, view performance metrics 289 ports, view performance metrics 289 ports, view performance metrics 289 inactive data 364 interfaces overview 43 inventory reports 475, 482, 483, 502
н	KB, See KiB key performance indicators for resources that run IBM Spectrum Virtualize
hardware zoning 82 hierarchies application 411 department 411	viewing <u>273</u> KiB <u>79</u>

L	P
license	PB, See PiB
capacity limit 30	performance
compliance	, storage systems, comparing 100
configuring <u>37</u>	buffer credits 108
model	controls for the performance view <u>266</u>
capacity <u>30</u> , <u>34</u> , <u>37</u>	exporting data <u>98</u>
enclosure <u>30</u> , <u>34</u> , <u>37</u>	hot spots <u>87</u>
modelconfiguring	key indicators 274
capacity 30, 34, 37	key performance indicators for resources that run IBN
enclosure 30, 34, 37	Spectrum Virtualize 273
license restrictions <u>26</u>	metrics
licensing	disk 277
examples 38	IBM Spectrum Accelerate 311
lists hiding columns 455	port <u>277</u> volume <u>2</u> 77
reorganizing columns 455	XIV 311
showing columns 455	monitoring 260
local server 77	performance view 265, 272
tocat server <u>rr</u>	troubleshooting 104
	viewing alerts 275
M	performance alert for applications
maintenance 60, 61	editing 176
managed disks	performance alert for general groups
capacity reports 481	editing 182
master server 77	performance alert for resource
max deviation percentage	editing 162
default range 446	performance data
MB, See MiB	analysis period, change <u>449</u>
metricsperformance reports 277	exporting <u>276</u>
metricsXIV 311	exporting performance data 276
MiB 79	performance metrics
modifying 423	Dell EMC storage systems 321
monitor	IBM Spectrum Scale 317
capacity and space usage 94	IBM Spectrum Virtualize 289
	NetApp storage systems 335
N	ONTAP 9 <u>335</u>
	Unity <u>321</u> VMAX 321
N_Port address identifier 82	VNX 321
N_Port_Name 82	volume cache (VC) 289
Netcool/OMNIbus alert notifications 187	volume copy cache (VCC) 289
Node_Name 82	performance metricsFlashArray//MFlashArray//X
non standard zones	Pure storage systems 339
fabric assist <u>84</u>	performance metricsHitachi VSP storage systems 332
protocol <u>84</u>	performance reports 489
quick loop 84	performance view
notifications 62, 63	accessing 265, 272
notifications for alerts 226	bookmarks 272
	controls 266
0	favorites 272
	saving URLs <u>272</u>
ONTAP 9	viewing the performance of internal resources <u>269</u>
disks, view performance metrics 335	viewing the performance of related resources <u>269</u>
nodes, view performance metrics 335	PiB <u>79</u>
pools, view performance metrics 335	pool)
ports, view performance metrics 335	capacity <u>240</u>
volumes, view performance metrics 335 Opening the management GUI for a resource 452	pools
overview 68	balance, analyze <u>445</u>
OVER VIEW OO	capacity view 358
	key storage values 387
	modifying properties <u>450</u> properties that can be modified 450
	DI ODEI LIES LIIAL CAII DE HIUUIHEU 430

pools (continued)	reports (continued)
tiering, analyze <u>443</u>	consumer reports, troubleshoot 500
port name zoning <u>82</u>	create <u>492</u>
port zoning 82	custom 473, 484, 492, 499
PortId zoning 82	custom reports, troubleshoot 500
predefined capacity reports 502	delete 499
predefined inventory reports 502	email 473
predefined reports	email reports, troubleshoot 500
edit 499	exporting 473
managed disks 481	inventory 483
pools 479	managed disks 481
pools 475 pools by tiers 481	performance 489
	•
volumes by servers 480	pools by tiers 481
volumes capacity report, create <u>478</u>	predefined 499
probes	predefined report, create 478
automated scheduling 140	predefined reports, save <u>475</u> , <u>479</u> , <u>482</u>
configuring alerts <u>142</u>	run now <u>499</u>
problems	schedule <u>488</u>
report <u>xi</u>	send by email <u>488</u>
product releases <u>60</u>	volumes by servers <u>480</u>
product updates 60, 61	resources
provisioned capacity)	adding 115
usage 244	adding fabrics for monitoring 129
provisioning	adding hypervisors for monitoring 132, 133
custom requirements for 424	adding rollup servers 136
overview 72, 420	adding storage systems for monitoring
shares 436, 437	Hitachi 121
volumes 435, 436	adding switches for monitoring 129
Pure FlashArray//M and FlashArray//X	adding to a capacity pool <u>426</u>
disks, view performance metrics 339	manually adding resources to applications 403
nodes, view performance metrics 339	modifying properties 450
pools, view performance metrics 339	properties that can be modified 450
ports, view performance metrics 339	removing <u>137</u>
volumes, view performance metrics 339	removing from a capacity pool 427
	restrictions
R	based on license <u>26</u>
IX	based on role <u>26</u>
reader feedback, sending xi	rfe 63
reclamation	role restrictions 26
block storage systems 446	roles 80
excluding 448	rollup mode
recommendations 447, 448	icon 77
	limitations and considerations 137
storage systems 448	master server 77
viewing volumes 447	subordinate server 77
volumes <u>446</u> , <u>448</u>	the control of the co
Reclamation	rollup server connections 77
storage <u>75</u>	rollup servers
volumes 75	adding 136
refresh packs 60	user roles <u>116</u>
remove, capacity limits 351	
removing 405, 410	S
removing applications 405	3
removing departments 410	SAN Volume Controller
removing resources 137	publications ix
reporting	script
dashboard 233	·
	performance data 277
reporting problems <u>xi</u>	security announcements 62
reports	security fixes 60, 61
application, performance <u>489</u>	security notifications <u>62</u> , <u>63</u>
capacity <u>473</u> , <u>483</u> , <u>484</u>	sending comments <u>xi</u>
capacity summary 498	server alerts
chargeback 492, 494	triggering conditions 219
chargeback reports, troubleshoot 500	servers
consumer 492, 496, 497	adding

servers (continued)	storage systems (continued)
adding (continued)	buffer credits 108
importing servers with agents from a file 135	exporting performance data 276
adding by deploying an agent <u>134</u>	horizontal bar charts, view key information 238
agentless, reviewing <u>101</u>	IBM Spectrum Scale, adding
modifying properties <u>450</u>	openstack access, configuring 129
optimized volumes, collocate 90	performance data collection, configuring 128
properties that can be modified 450	probe, configuring 128
tiering, analyze 442	sudo access, configuring 127
viewing information in the GUI 231	modifying properties 450
viewing status in the web-based GUI 247	performance data, exporting 276
service xi	performance, comparing 100
service class	properties that can be modified 450
deleting 424	user roles 116
satisfying custom requirements of a 424	viewing enclosures 244
viewing properties of 423	viewing information in the GUI 231
viewing storage constraints of 423	viewing status in the web-based GUI 247
viewing users who can provision by using 423	storage virtualizers
viewing volumes or shares associated with 423	tiering, analyze 442
service management connect xi	volumes, analyze placement 87
shares 150	volumes, optimize 87
shortfall 366	subcomponents
SMC xi	creating subcomponents for existing applications 403
SNMP alert notifications 187	
social media	subdepartments
	creating 408
blogs <u>17, 19</u>	subordinate servers
videos 17	adding 77
software zoning 82	limitations and considerations 137
sorting rows in lists 455	removing 77
space usage	subscribe <u>62, 63</u>
block storage systems <u>348</u>	suggestions <u>63</u>
containers 348	support <u>xi</u>
file storage systems 348	Support
file system pools 348	contact <u>xi</u>
file systems <u>348</u>	switch alerts
filesets <u>348</u>	triggering conditions <u>215</u>
monitoring <u>348</u>	switches
object storage systems <u>348</u>	adding <u>129,</u> <u>132</u>
pools <u>348</u>	agent types for monitoring 132
tiers <u>348</u>	modifying properties <u>450</u>
volumes <u>348</u>	properties that can be modified 450
Sponsor user program	viewing information in the GUI 231
advantages <u>xxxii</u>	viewing status in the web-based GUI 247
overview <u>xxxii</u>	
standard provisioning 381	T
starting	1
IBM Spectrum Control web-based GUI 44	tasks
status of resources	analysis execution 463, 464
viewing the status of resources 253	creating provisioning tasks in the web-based GUI 460
status of resources in the web-based GUI	creating tiering storage, balancing pools, and
acknowledging the status of a resource 254	transforming-storage tasks in the GUI 462
acknowledging the status of a top-level resource 254	deleting in the web-based GUI 466
acknowledging the status of an internal resource 255	exporting provisioning tasks in the web-based GUI 461
viewing the overall status of resources 252	implementing optimization recommendations in
storage pools	
tier level, setting 441	theweb-based GUI 463
tiering, analyze 443	managing in the GUI 455
storage system alerts	managing provisioning tasks in the GUI 459
triggering conditions 188	managing tiering storage, balancing pools, and
storage system internal resource alerts	transforming-storage tasks in the GUI 461
triggering conditions 194	pausing, resuming, and canceling analysis-execution
	tasks in the web-based GUI 465
storage systems	running provisioning tasks in the web-based GUI <u>460</u>
adding	running tiering analysis, transform-plan, and balance-
Hitachi <u>121</u>	analysis tasks in the web-based GUI 462

tasks (continued)	troubleshooting (continued)
scheduling the implementation of optimization	pools <u>348</u>
recommendations in theweb-based GUI 464	tiers <u>348</u>
setting the number of task runs displayed 459	volumes <u>348</u>
viewing all tasks in the web-based GUI <u>456</u>	tutorials
viewing details in the web-based GUI <u>457</u>	connectivity information 97
viewing in the GUI <u>456</u>	exporting performance data <u>98</u>
viewing pending tasks <u>456</u>	NPIV connections 97
viewing task logs in the web-based GUI 458	reviewing agentless servers 101
viewing tasks for a resource in the web-based GUI <u>457</u>	storage systems, comparing performance 100
viewing tasks for a resource type in the web-based GUI	troubleshooting application performance 104
<u>457</u>	
TB, See TiB	U
temporary zone sets <u>82</u>	
thin-provisioning	units of measurement 79
data reduction savings <u>381</u>	upgrading
effective capacity <u>381</u>	configuring licensing 37
thin-provisioning savings <u>381</u>	user roles
total capacity savings <u>381</u>	hypervisors 116
TiB 79	rollup servers 116
tier levels	storage systems 116
storage pools, setting <u>441</u>	VMware vSphere data sources 116
tiering	users
analysis period <u>449</u>	assigning roles 80
I/O density <u>443</u>	using
I/O rate <u>443</u>	REST API 85, 468
pools <u>440</u>	_ _ _
pools, analyze <u>440</u>	V
pools, assign tier levels <u>440</u>	V
tiering thresholds	vCenter Server system
I/O rate, modify <u>444</u>	adding 132, 133
tiers	videos 17
capacity reports <u>481</u>	viewing
renaming <u>442</u>	connectivity information 97
TPCD.cfg	NPIV connections 97
saveNonRoot parameter, add <u>150</u>	viewing details of a department 410
trademarks 519	viewing details of an application 404
translations	Viewing
browser locale requirement <u>x</u>	the overall status of a storage environment 233
triggering actions 226	viewing application details 404
triggering conditions	viewing department details 410
alerts 188	viewing enclosures 244
computer alerts 219	viewing the status of resources 253
fabric alerts 218	volume collocation 90
hypervisor alerts 210	volumes
server alerts 219	compressed, view capacity charts 348
storage system alerts <u>188</u>	thin-provisioned, view capacity charts 348
storage system internal resource alerts 194	tiering 88
switch alerts <u>215</u>	tiering by I/O density 443
troubleshoot	tiering by I/O rate 443
chargeback reports 500	tiering, analyze 443
consumer reports 500	workload 445
custom reports 500	volumes, shares 434
email reports <u>500</u>	volumes)
troubleshooting	provisioned capacity usage 244
block storage systems 348	VSANs 83
capacity 348	<u></u>
containers 348	147
file storage systems 348	W
file system pools 348	web-based GUI
file systems 348	
filesets 348	acknowledging the status of a resource 254 acknowledging the status of a top-level resource 254
object storage systems 348	
performance <u>260</u>	acknowledging the status of an internal resource <u>255</u>

```
web-based GUI (continued)
    alert severities 158
    data collection
        creating performance monitors 144
        modifying performance monitors 146
        verifying that a performance monitor is running 145
        verifying that a probe is running 140
    filtering 454
    Monitoring the status of resources 247
    status of resources 247
        analysis execution 463, 464
        creating provisioning tasks 460
        creating tiering storage, balancing pools, and
        transforming-storage tasks 462
        deleting 466
        exporting provisioning tasks 461
        implementing optimization recommendations 463
        pausing, resuming, and canceling analysis-
        execution tasks 465
        running provisioning tasks 460
        running tiering analysis, transform plan, and
        balance-analysis tasks 462
        scheduling the implementation of optimization
        recommendations 464
    viewing all tasks 456
    viewing information about top-level resources 236
    viewing performance monitor logs 149
    viewing performance monitors 148
    viewing performance monitors for a resource 148
    viewing task details 457
    viewing task for a resource 457
    viewing task for a resource type 457
    viewing task logs 458
    viewing the condition of top-level resources 252
    viewing the overall status of resources 252
Windows 509, 510
workload
    volumes 445
WWN zoning 82
X
XIVperformance metrics 311
Z
zone
    concepts 82
    default 82
    hardware 82
    membership 82
    non standard 84
    port 82
    software 82
    supported switches 84
zone set
    temporary 82
zones 81
zoning
    automatic 82
    during provisioning 82
```

IBW.

Product Number: 5725-F93, 5725-G33

SC28-3133-00

