IBM Tivoli Storage Manager
Version 7.1.3

Multisite Disk Solution Guide

IBM

IBM Tivoli Storage Manager
Version 7.1.3

Multisite Disk Solution Guide

IBM

Note:
Before you use this information and the product it supports, read the information in "Notices" on page 157.

Contents

About this publication

This publication provides information about planning for, implementing, monitoring, and operating a data protection solution that uses IBM® Tivoli® Storage Manager best practices.

Who should read this guide

This guide is intended for anyone who is registered as an administrator for Tivoli Storage Manager. A single administrator can manage Tivoli Storage Manager, or several people can share administrative responsibilities.

You should be familiar with the operating system on which the server resides and the communication protocols required for the client/server environment. You also need to understand the storage management practices of your organization, such as how you are currently backing up workstation files and how you are using storage devices.

Publications

The Tivoli Storage Manager product family includes IBM Tivoli Storage FlashCopy® Manager, IBM Tivoli Storage Manager for Space Management, IBM Tivoli Storage Manager for Databases, and several other storage management products from IBM.

To view IBM product documentation, see http://www.ibm.com/support/knowledgecenter.

New for Tivoli Storage Manager Version 7.1.3

Many features in IBM Tivoli Storage Manager Version 7.1.3 are new for Tivoli Storage Manager users. These new features are designed to optimize data storage and improve product usability.

For a complete list of new features in this release, see What's new (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.common/r_wn_tsmserver.html).

Part 1. Planning for a multisite disk data protection solution

Plan for a multisite disk data protection solution with IBM Tivoli Storage Manager servers at two sites that use data deduplication and replication.

Implementation methods

You can configure Tivoli Storage Manager servers for a multisite disk solution in the following ways:

Configure servers by using the Operations Center and administrative commands
This documentation provides steps to configure a range of storage systems and the server software for your solution. Configuration tasks are completed by using wizards and options in the Operations Center and Tivoli Storage Manager commands. For information about getting started, see the "Planning roadmap."

Configure the servers by using automated scripts
For detailed guidance on configuration with specific IBM Storwize® storage systems, and by using automated scripts to configure each Tivoli Storage Manager server, see the Tivoli Storage Manager blueprints. The documentation and scripts are available on IBM developerWorks® at: IBM Spectrum Protect™ Blueprints (https://www.ibm.com/developerworks/community/wikis/home/wiki/Tivoli%20Storage%20Manager/page/IBM%20Spectrum%20Protect%20Blueprints).

The blueprint documentation does not include steps for installing and configuring the Operations Center, or setting up secure communications by using Transport Security Layer (TLS). Replication is configured by using commands after each server is set up. An option for using Elastic Storage Server, based on IBM Spectrum Scale™ technology, is included.

Planning roadmap

Plan for a Tivoli Storage Manager multisite disk solution by reviewing the architecture layout in Figure 1 on page 2 and then completing the roadmap tasks that follow the diagram.
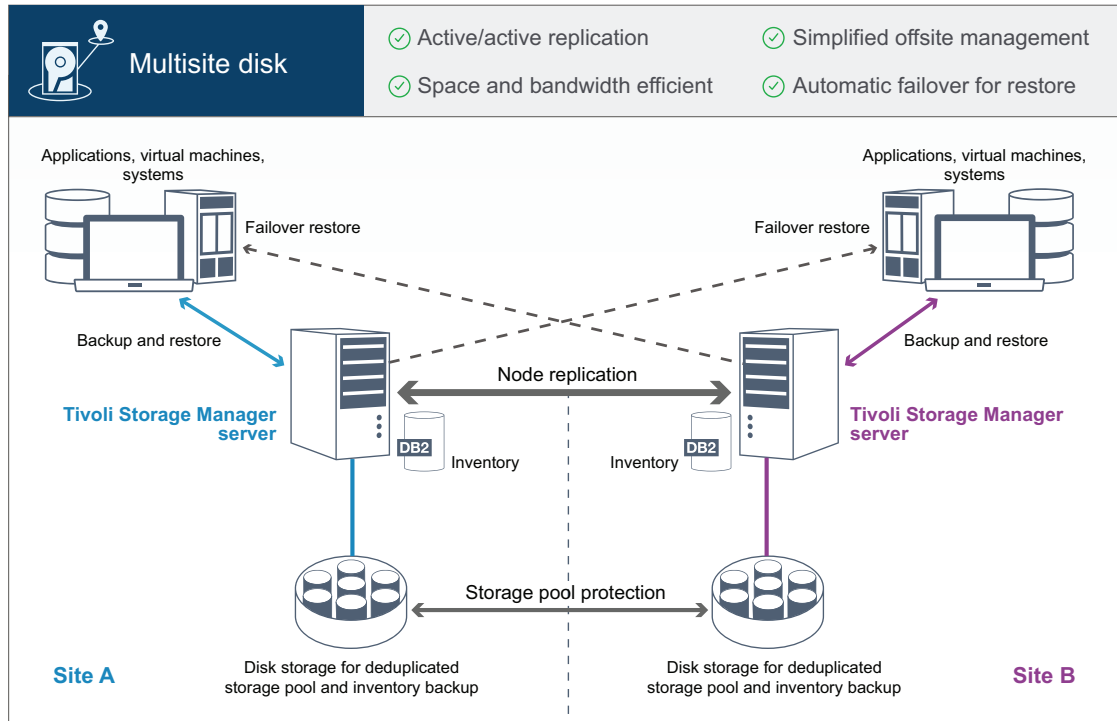
Figure 1. Multisite disk solution

The following steps are required to plan properly for a multisite disk environment.

1. Select your system size.
2. Plan for the sites.
3. Meet system requirements for hardware and software.
4. Record values for your system configuration in the planning worksheets.
5. Plan for storage.
6. Plan for security.
   a. Plan for administrator roles.
   b. Plan for secure communications.
   c. Plan for storage of encrypted data.
   d. Plan for firewall access.

Chapter 1. Selecting a system size

Select the size of the Tivoli Storage Manager server based on the amount of data that you manage and the systems to be protected.

About this task

You can use the information in the table to determine the size of Tivoli Storage Manager server that is required, based on the amount of data that you manage.

The following table describes the volume of data that a server manages. This amount includes all versions. The daily amount of data is how much new data you back up each day. Both the total managed data and daily amount of new data are measured as the size before any data reduction.

Table 1. Determining the size of the Tivoli Storage Manager server

Total managed data	Daily amount of new data to back up	Required server size
45 TB - 180 TB	Up to 6 TB per day	Small
200 TB - 800 TB	6 - 20 TB per day	Medium
1000 TB - 4000 TB	20 - 100 TB per day	Large

The daily backup values in the table are based on test results with 128 MB sized objects, which are used by Tivoli Storage Manager for Virtual Environments. Workloads that consist of objects that are smaller than 128 KB might not be able to achieve these daily limits.

Chapter 2. Planning the sites

Review use cases and evaluate the factors to provide the most efficient data protection for the multisite disk solution for Tivoli Storage Manager.

Use cases

The multisite disk solution creates at least one copy of backed-up data. If the Tivoli Storage Manager servers are at separate locations, the backed-up replica is maintained offsite. Although your company might benefit from a multisite disk solution for various reasons, the most common reasons to use a multisite disk solution include the following replication scenarios:

Replication from the primary site to the disaster recovery site

> In this scenario, data that is backed up from the primary site, Site A, is replicated to a Tivoli Storage Manager server at the secondary, disaster recovery site, Site B. If a disaster occurs at Site A, such as failure of the server, you can use the Tivoli Storage Manager server at Site B to recover systems. Alternatively, you can use the server at Site A to restore primary storage pool data at Site B, such as after a disk storage failure at Site B..

Mutual replication at two active sites

> In this scenario, local data at each site is backed up by the Tivoli Storage Manager servers at both Site A and Site B. Data that is backed up from Site A is replicated to Site B, and backed-up data from Site B is replicated to Site A. If data that was backed up is lost at Site A, you can use the server at Site B to recover storage pool data to the server at Site A. If Site A is no longer available, you can recover the replicated data for Site A to a new system at Site B. You must size the server resources to ensure that either server has sufficient capacity to back up and restore all client nodes as part of your disaster recovery plan.

Protect remote Tivoli Storage Manager servers to the primary site

> In this scenario, you configure remote Tivoli Storage Manager servers that are relatively small to replicate data that is backed up to a larger Tivoli Storage Manager server at the primary site. If bandwidth is limited, it might not be practical to restore systems to the remote sites. In this case, you might want to recover systems at the primary site before you replicate the backed-up data to the remote servers.

Factors to evaluate

Before you implement a multisite disk solution, evaluate the following factors:

Network bandwidth

> The network must have sufficient bandwidth for the expected data transfers between nodes, for replication, and for the cross-site restore operations that are required for disaster recovery. Before you proceed with testing replication throughput, ensure that your network can handle the replication traffic. Calculate the required network bandwidth for the steady-state requirement by applying the guidelines in Estimating network

bandwidth required for replication (V7.1.1) (http://www.ibm.com/
support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.srv.doc/
t_repl_est_bw.html).

The network connection is often a shared resource. Plan the time of day to
schedule node replication to run to avoid a conflict with other resource
users. Also, network controls might limit Tivoli Storage Manager activity to
only a portion of the bandwidth. There are no controls in Tivoli Storage
Manager to restrict network usage.

Resources for the initial replication

To set up the data protection solution across two sites, you must replicate
data initially from Site A to the target server at Site B. To ensure that the
initial replication is successful, you must determine whether you have the
network bandwidth, processor resources, and time available to replicate the
data. You might have to plan for replicating the initial full backups across
several days. If you cannot extend the schedule for the initial backups, you
can replicate data from Site A to Site B without using the network. For
example, you can export and import the backed-up data by using media or
you can temporarily locate the source and target servers on the same site.

Daily data ingestion

For the multisite disk solution, the daily data ingestion and total data
retention must be within the capacity of the Tivoli Storage Manager
configurations. For example, a large Tivoli Storage Manager configuration
has a data ingestion capacity of up to 100 TB per day, including node
replication. In cases where the backup requirements exceed the capacity of
a single Tivoli Storage Manager server, you can configure a solution that
uses multiple Tivoli Storage Manager servers to achieve the required
capacity.

Server configuration

The Tivoli Storage Manager server configuration must meet or exceed the
requirements for the multisite disk solution.

Single replica of backed-up data

The multisite disk solution is most efficient when a single, offsite copy of
the backed-up data meets your data protection and risk mitigation
requirements. In this case, the single copy of the data is maintained off-site
at the location of a Tivoli Storage Manager replication server.

Related reference:

Chapter 3. System requirements for a multisite disk solution

After you select the solution that best fits your data protection requirements, review the system requirements to plan for implementation of the data protection solution.

Ensure that your system meets the hardware and software prerequisites for the size of server that you plan to use.

Related information:

▶ Tivoli Storage Manager Supported Operating Systems (http://www.ibm.com/support/docview.wss?uid=swg21243309)

Hardware requirements

Hardware requirements are based on system size. Choose equivalent or better components than those items that are listed to ensure optimum performance for your environment.

For a definition of system sizes, see Chapter 1, "Selecting a system size," on page 3.

The following table includes minimum hardware requirements for the server and storage, based on the size of the server that you plan to build. If you are using local partitions (LPARs) or work partitions (WPARs), adjust the network requirements to take account of the partition sizes.

Hardware requirements for the Operations Center are included in the table, except for the database and archive log space that the Operations Center uses to hold records for managed clients. This value is based on the number of clients that are monitored on a server. For guidance about how to calculate this estimate, see Estimating database space requirements for the Operations Center.

If you do not plan to install the Operations Center on the same system as the server, you can estimate system requirements separately. For details about calculating system requirements for the Operations Center, see Tivoli Storage Manager Operations Center system requirements calculator (http://www.ibm.com/support/docview.wss?uid=swg21641684).

Hardware component	Small system	Medium system	Large system
Server processor	AIX 6 processor cores, 3.42 GHz or faster Linux Windows 12 processor cores, 1.9 GHz or faster	AIX 8 processor cores, 3.42 GHz or faster Linux Windows 16 processor cores, 2.0 GHz or faster	AIX 20 processor cores, 3.42 GHz Linux Windows 32 processor cores, 2.0 GHz or faster
Server memory	64 GB RAM	128 GB RAM	192 GB RAM
Network	• 10 GB Ethernet (1 port) • 8 GB Fibre Channel adapter (2 ports)	• 10 GB Ethernet (2 ports) • 8 GB Fibre Channel adapter (2 ports)	• 10 GB Ethernet (4 ports) • 8 GB Fibre Channel adapter (4 ports)

Hardware component	Small system	Medium system	Large system
Storage	• 1.3 TB inventory[1] • 46 TB deduplicated directory-container storage pool	• 2 TB inventory[1] • 200 TB deduplicated directory-container storage pool	• 4 TB inventory[1] • 1000 TB deduplicated directory-container storage pool
[1]Plus extra space that is required to hold records that are managed by the Operations Center.			

Estimating database space requirements for the Operations Center

Managing the Operations Center on a Tivoli Storage Manager server represents an extra workload. This workload requires extra space for database operations. Review the following guidelines to help estimate how much space your server requires.

Database space
> The Operations Center uses approximately 1.2 GB of database space for every 1000 clients that are monitored on a server. For example, consider a hub server with 2000 clients that also manages three spoke servers, each with 1500 clients. This configuration has a total of 6500 clients across the four Tivoli Storage Manager servers and requires approximately 8.4 GB of database space. This value is calculated by rounding the 6500 clients up to the next closest 1000, which is 7000:
>
> ```
> 7 x 1.2 GB=8.4 GB
> ```

Archive log space
> The Operations Center uses approximately 8 GB of archive log space every 24 hours, for every 1000 clients. In the example of 6500 clients across the hub server and the spoke servers, 56 GB of archive log space is used over a 24-hour period for the hub server.
>
> For a given spoke server in the example, the archive log space that is used over 24 hours would be 16 GB. These estimates are based on the default status collection interval of 5 minutes. If you lower the collection interval from once every 5 minutes to once every 3 minutes, the space requirements increase. The following examples show the approximate increase in the log space requirement with a collection interval of once every 3 minutes:
>
> • Hub server: 56 GB to approximately 94 GB
> • Each spoke server: 16 GB to approximately 28 GB
>
> Increase the archive log space so that you have sufficient space available to support the Operations Center, without affecting the existing server operations.

Hardware requirements for the second server

If you are planning to set up your sites so that everything at the first site is replicated to the second site, hardware requirements will be identical at both sites. If you want to only replicate a subset of data to your second site, storage and network requirements might be reduced.

Software requirements

Documentation for the multisite disk solution includes installation and configuration tasks for the following operating systems. You must meet the minimum software requirements that are listed.

AIX® systems

Type of software	Minimum software requirements
Operating system	IBM AIX 7.1
Gunzip utility	The gunzip utility must be available on your system before you install or upgrade the Tivoli Storage Manager Version 7 server. Ensure that the gunzip utility is installed and the path to it is set in the PATH environment variable.
File system type	JFS2 file systems AIX systems can cache a large amount of file system data, which can reduce memory that is required for Tivoli Storage Manager server and DB2® processes. To avoid paging with the AIX server, use the rbrw mount option for the JFS2 file system. Less memory is used for the file system cache and more is available for Tivoli Storage Manager. Do not use the file system mount options, Concurrent I/O (CIO), and Direct I/O (DIO), for file systems that contain the Tivoli Storage Manager database, logs, or storage pool volumes. These options can cause performance degradation of many server operations. Tivoli Storage Manager and DB2 can still use DIO where it is beneficial to do so, but Tivoli Storage Manager does not require the mount options to selectively take advantage of these techniques.
Other software	Korn Shell (ksh)

Linux systems

Type of software	Minimum software requirements
Operating system	Red Hat Enterprise Linux 7 (x86_64)
Libraries	GNU C libraries, Version 2.3.3-98.38 or later that is installed on the Tivoli Storage Manager system. Red Hat Enterprise Linux Servers: • `libaio` • `libstdc++.so.6` (32-bit and 64-bit packages are required) • `numactl.x86_64`
File system type	Format database-related file systems with ext3 or ext4. For storage pool-related file systems, use XFS.
Other software	Korn Shell (ksh)

Windows systems

Type of software	Minimum software requirements
Operating system	Microsoft Windows 2012 R2 (64-bit)

Type of software	Minimum software requirements
File system type	NTFS
Other software	Windows 2012 R2 with .NET Framework 4.5 is installed and enabled. The following User Account Control policies must be disabled: • User Account Control: Admin Approval Mode for the Built-in Administrator account • User Account Control: Run all administrators in Admin Approval Mode

Related tasks:

➥ Setting AIX network options (V7.1.1) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.perf.doc/t_network_aix_srv_clnt.html)

Chapter 4. Planning worksheets

Use the planning worksheets to record values that you use to set up your system and configure the Tivoli Storage Manager server. Use the best practice default values that are listed in the worksheets.

Each worksheet helps you prepare for different parts of the system configuration by using best practice values:

Server system preconfiguration
> Use the preconfiguration worksheets to plan for the file systems and directories that you create when you configure file systems for Tivoli Storage Manager during system setup. All directories that you create for the Tivoli Storage Manager server must be empty.

Tivoli Storage Manager configuration
> Use the Tivoli Storage Manager configuration worksheets when you configure the Tivoli Storage Manager server. Default values are recommended for most items, except where noted.

AIX

Table 2. Worksheet for preconfiguration of an AIX server system

Item	Default value	Your value	Minimum directory size	Notes
TCP/IP port address for communications with the Tivoli Storage Manager server	1500		Not applicable	Ensure that this port is available when you install and configure the operating system The port number can be a number in the range 1024 - 32767.
Directory for the server instance	/home/tsminst1/tsminst1		50 GB	If you change the value for the server instance directory from the default, also modify the DB2 instance owner value in Table 3 on page 13.
Directory for server installation	/		Available space that is required for the directory: 5 GB	
Directory for server installation	/usr		Available space that is required for the directory: 5 GB	
Directory for server installation	/var		Available space that is required for the directory: 5 GB	

Table 2. Worksheet for preconfiguration of an AIX server system (continued)

Item	Default value	Your value	Minimum directory size	Notes
Directory for server installation	`/tmp`		Available space that is required for the directory: 5 GB	
Directory for server installation	`/opt`		Available space that is required for the directory: 10 GB	
Directory for the active log	`/tsminst1/TSMalog`		• Small and medium: 140 GB • Large: 300 GB	When you create the active log during the initial configuration of the server, set the size to 128 GB.
Directory for the archive log	`/tsminst1/TSMarchlog`		• Small: 1 TB • Medium: 3 TB • Large: 4 TB	

Directories for the database	`/tsminst1/TSMdbspace00`
	`/tsminst1/TSMdbspace01`
	`/tsminst1/TSMdbspace02`
	`/tsminst1/TSMdbspace03`
	...
			Minimum total space for all directories:
			• Small: At least 1 TB
			• Medium: At least 2 TB
			• Large: At least 4 TB
				Create a minimum number of file systems for the database, depending on the size of your Tivoli Storage Manager system:
				• Small: At least 4 file systems
				• Medium: At least 8 file systems
				• Large: At least 8 file systems

Directories for storage	`/tsminst1/TSMfile00`
	`/tsminst1/TSMfile01`
	`/tsminst1/TSMfile02`
	`/tsminst1/TSMfile03`
	...
			Minimum total space for all directories:
			• Small: At least 38 TB
			• Medium: At least 180 TB
			• Large: At least 980 TB
				Create a minimum number of file systems for storage, depending on the size of your Tivoli Storage Manager system:
				• Small: At least 10 file systems
				• Medium: At least 20 file systems
				• Large: At least 40 file systems

Table 2. Worksheet for preconfiguration of an AIX server system (continued)

Item	Default value	Your value	Minimum directory size	Notes
Directories for database backup	/tsminst1/TSMbkup00 /tsminst1/TSMbkup01 /tsminst1/TSMbkup02 /tsminst1/TSMbkup03		Minimum total space for all directories: • Small: At least 3 TB • Medium: At least 10 TB • Large: At least 16 TB	Create a minimum number of file systems for backing up the database, depending on the size of your Tivoli Storage Manager system: • Small: At least 2 file systems • Medium: At least 4 file systems • Large: At least 4 file systems

Table 3. Worksheet for Tivoli Storage Manager configuration

Item	Default value	Your value	Notes
DB2 instance owner	tsminst1		If you changed the value for the server instance directory in Table 2 on page 11 from the default, also modify the value for the DB2 instance owner.
DB2 instance owner password	passw0rd		Select a different value for the instance owner password than the default. Ensure that you record this value in a secure location.
Primary group for the DB2 instance owner	tsmsrvrs		
Server name	The default value for the server name is the system host name.		
Server password	passw0rd		Select a different value for the server password than the default. Ensure that you record this value in a secure location.
Tivoli Storage Manager administrator ID (user ID for the Tivoli Storage Manager server instance)	admin		
Tivoli Storage Manager administrator ID password	passw0rd		Select a different value for the administrator password than the default. Ensure that you record this value in a secure location.

Table 3. Worksheet for Tivoli Storage Manager configuration (continued)

Item	Default value	Your value	Notes
Schedule start time	`22:00`		The default schedule start time begins the client workload phase, which is predominantly the client backup and archive activities. During the client workload phase, server resources support client operations. Normally, these operations are completed during the nightly schedule window.
			Schedules for server maintenance operations are defined to begin 10 hours after the start of the client backup window.

Linux

Table 4. Worksheet for preconfiguration of a Linux server system

Item	Default value	Your value	Minimum directory size	Notes
TCP/IP port address for communications with the Tivoli Storage Manager server	1500		Not applicable	Ensure that this port is available when you install and configure the operating system The port number can be a number in the range 1024 - 32767.
Directory for the server instance	`/home/tsminst1/tsminst1`		25 GB	If you change the value for the server instance directory from the default, also modify the DB2 instance owner value in Table 5 on page 15.
Directory for the active log	`/tsminst1/TSMalog`		• Small and medium: 140 GB • Large: 300 GB	
Directory for the archive log	`/tsminst1/TSMarchlog`		• Small: 1 TB • Medium: 3 TB • Large: 4 TB	

Table 4. Worksheet for preconfiguration of a Linux server system (continued)

Item	Default value	Your value	Minimum directory size	Notes
Directories for the database	/tsminst1/TSMdbspace00 /tsminst1/TSMdbspace01 /tsminst1/TSMdbspace02 /tsminst1/TSMdbspace03 ...		Minimum total space for all directories: • Small: At least 1 TB • Medium: At least 2 TB • Large: At least 4 TB	Create a minimum number of file systems for the database, depending on the size of your Tivoli Storage Manager system: • Small: At least 4 file systems • Medium: At least 8 file systems • Large: At least 8 file systems
Directories for storage	/tsminst1/TSMfile00 /tsminst1/TSMfile01 /tsminst1/TSMfile02 /tsminst1/TSMfile03 ...		Minimum total space for all directories: • Small: At least 38 TB • Medium: At least 180 TB • Large: At least 980 TB	Create a minimum number of file systems for storage, depending on the size of your Tivoli Storage Manager system: • Small: At least 10 file systems • Medium: At least 20 file systems • Large: At least 40 file systems
Directories for database backup	/tsminst1/TSMbkup00 /tsminst1/TSMbkup01 /tsminst1/TSMbkup02 /tsminst1/TSMbkup03		Minimum total space for all directories: • Small: At least 3 TB • Medium: At least 10 TB • Large: At least 16 TB	Create a minimum number of file systems for backing up the database, depending on the size of your Tivoli Storage Manager system: • Small: At least 2 file systems • Medium: At least 4 file systems • Large: At least 4 file systems

Table 5. Worksheet for Tivoli Storage Manager configuration

Item	Default value	Your value	Notes
DB2 instance owner	tsminst1		If you changed the value for the server instance directory in Table 4 on page 14 from the default, also modify the value for the DB2 instance owner.

Table 5. Worksheet for Tivoli Storage Manager configuration (continued)

Item	Default value	Your value	Notes
DB2 instance owner password	`passw0rd`		Select a different value for the instance owner password than the default. Ensure that you record this value in a secure location.
Primary group for the DB2 instance owner	`tsmsrvrs`		
Server name	The default value for the server name is the system host name.		
Server password	`passw0rd`		Select a different value for the server password than the default. Ensure that you record this value in a secure location.
Tivoli Storage Manager administrator ID (user ID for the Tivoli Storage Manager server instance)	`admin`		
Tivoli Storage Manager administrator ID password	`passw0rd`		Select a different value for the administrator password than the default. Ensure that you record this value in a secure location.
Schedule start time	`22:00`		The default schedule start time begins the client workload phase, which is predominantly the client backup and archive activities. During the client workload phase, server resources support client operations. Normally, these operations are completed during the nightly schedule window. Schedules for server maintenance operations are defined to begin 10 hours after the start of the client backup window.

Windows

Because many volumes are created for the Tivoli Storage Manager server, configure the server by using the Windows feature of mapping disk volumes to directories rather than to drive letters.

For example, `C:\tsminst1\TSMdbpsace00` is a mount point to a volume with its own space. It is mapped to a directory under the C: drive, but does not take up space from the C: drive. The exception is the server instance directory, `C:\tsminst1`, which can be a mount point or a regular directory.

Table 6. Worksheet for preconfiguration of a Windows server system

Item	Default value	Your value	Minimum directory size	Notes
TCP/IP port address for communications with the Tivoli Storage Manager server	1500		Not applicable	Ensure that this port is available when you install and configure the operating system The port number can be a number in the range 1024 - 32767.
Directory for the server instance	`C:\tsminst1`		25 GB	If you change the value for the server instance directory from the default, also modify the DB2 instance owner value in Table 7 on page 18.
Directory for the active log	`C:\tsminst1\TSMalog`		• Small and medium: 140 GB • Large: 300 GB	
Directory for the archive log	`C:\tsminst1\TSMarchlog`		• Small: 1 TB • Medium: 3 TB • Large: 4 TB	

Directories for the database	`C:\tsminst1\TSMdbspace00`
	`C:\tsminst1\TSMdbspace01`
	`C:\tsminst1\TSMdbspace02`
	`C:\tsminst1\TSMdbspace03`
	...
			Minimum total space for all directories:
			• Small: At least 1 TB
			• Medium: At least 2 TB
			• Large: At least 4 TB
				Create a minimum number of file systems for the database, depending on the size of your Tivoli Storage Manager system:
				• Small: At least 4 file systems
				• Medium: At least 8 file systems
				• Large: At least 8 file systems

Table 6. Worksheet for preconfiguration of a Windows server system (continued)

Item	Default value	Your value	Minimum directory size	Notes
Directories for storage	C:\tsminst1\TSMfile00 C:\tsminst1\TSMfile01 C:\tsminst1\TSMfile02 C:\tsminst1\TSMfile03 ...		Minimum total space for all directories: • Small: At least 38 TB • Medium: At least 180 TB • Large: At least 980 TB	Create a minimum number of file systems for storage, depending on the size of your Tivoli Storage Manager system: • Small: At least 10 file systems • Medium: At least 20 file systems • Large: At least 40 file systems
Directories for database backup	C:\tsminst1\TSMbkup00 C:\tsminst1\TSMbkup01 C:\tsminst1\TSMbkup02 C:\tsminst1\TSMbkup03		Minimum total space for all directories: • Small: At least 3 TB • Medium: At least 10 TB • Large: At least 16 TB	Create a minimum number of file systems for backing up the database, depending on the size of your Tivoli Storage Manager system: • Small: At least 2 file systems • Medium: At least 4 file systems • Large: At least 4 file systems

Table 7. Worksheet for Tivoli Storage Manager configuration

Item	Default value	Your value	Notes
DB2 instance owner	tsminst1		If you changed the value for the server instance directory in Table 6 on page 17 from the default, also modify the value for the DB2 instance owner.
DB2 instance owner password	pAssW0rd		Select a different value for the instance owner password than the default. Ensure that you record this value in a secure location.
Server name	The default value for the server name is the system host name.		
Server password	passw0rd		Select a different value for the server password than the default. Ensure that you record this value in a secure location.

Table 7. Worksheet for Tivoli Storage Manager configuration (continued)

Item	Default value	Your value	Notes
Tivoli Storage Manager administrator ID (user ID for the Tivoli Storage Manager server instance)	`admin`		
Tivoli Storage Manager administrator ID password	`passw0rd`		Select a different value for the administrator password than the default. Ensure that you record this value in a secure location.
Schedule start time	`22:00`		The default schedule start time begins the client workload phase, which is predominantly the client backup and archive activities. During the client workload phase, server resources support client operations. Normally, these operations are completed during the nightly schedule window. Schedules for server maintenance operations are defined to begin 10 hours after the start of the client backup window.

Chapter 5. Planning for storage

Choose the most effective storage technology for Tivoli Storage Manager components to ensure efficient server performance and operations.

Storage hardware devices have different capacity and performance characteristics, which determine how they can be used effectively with Tivoli Storage Manager. For general guidance on selecting the right storage hardware and set up for your solution, review the following guidelines.

Database and active log
- Use a fast disk for the Tivoli Storage Manager database and active log, for example with the following characteristics:
  - High performance, 15k rpm disk with Fibre Channel or serial-attached SCSI (SAS) interface
  - Solid-state disk (SSD)
- Isolate the active log from the database unless you use SSD or flash hardware
- When you create arrays for the database, use RAID level 5

Storage pool
- You can use less expensive and slower disks for the storage pool
- The storage pool can share disks for the archive log and database backup storage
- Use RAID level 6 for storage pool arrays to add protection against double drive failures when you use large disk types

Planning the storage arrays

Prepare for disk storage configuration by planning for RAID arrays and volumes, according to your system size.

You design storage arrays with size and performance characteristics that are suitable for one of the Tivoli Storage Manager server storage components, such as the server database or a storage pool. The storage planning activity must take account of drive type, RAID level, number of drives, the number of spare drives, and so on. In the solution configurations, storage groups contain internal-storage RAID arrays and consist of multiple physical disks that are presented as logical volumes to the system. When you configure the disk storage system, you create storage groups, or data storage pools, and then create storage arrays in the groups.

You create volumes, or LUNs, from the storage groups. The storage group defines which disks provide the storage that makes up the volume. When you create volumes, make them fully allocated. Faster disks types are used to hold the database volumes and active log volumes. Slower disk types can be used for the storage pool volumes, archive log, and database backup volumes.

Table 8 on page 22 and Table 9 on page 22 describe the layout requirements for storage groups and volume configuration.

Table 8. Components of storage group configuration

Component	Details
Server storage requirement	How the storage is used by the Tivoli Storage Manager server.
Disk type	Size and speed for the disk type that is used for the storage requirement.
Disk quantity	Number of each disk type that is needed for the storage requirement.
Hot spare capacity	Number of disks that are reserved as spares to take over if disk failures occur.
RAID level	Level of RAID array that is used for logical storage. The RAID level defines the type of redundancy that is provided by the array, for example 5 or 6.
RAID array quantity	Number of RAID arrays to be created.
DDMs per RAID array	How many disk drive modules (DDMs) are to be used in each of the RAID arrays.
Usable size per RAID array	Size that is available for data storage in each RAID array after accounting for space that is lost due to redundancy.
Total usable size	Total size that is available for data storage in the RAID arrays: [Quantity x Usable size].
Suggested storage group and array names	Preferred name to use for MDisks and MDisk groups.
Usage	Tivoli Storage Manager server component that uses part of the physical disk.

Table 9. Components of volume configuration

Component	Details
Server storage requirement	Requirement for which the physical disk is used.
Volume name	Unique name that is given to a specific volume.
Storage group	Name of the storage group from which the space is obtained to create the volume.
Size	Size of each volume.
Intended server mount point	Directory on the Tivoli Storage Manager server system where the volume is mounted.
Quantity	Number of volumes to create for a specific requirement. Use the same naming standard for each volume that is created for the same requirement.
Usage	Tivoli Storage Manager server component that uses part of the physical disk.

Examples

Configuration examples for storage groups and volumes are available at the following link: Example worksheets for planning storage arrays. The examples show how to plan the storage for different server sizes. In the example configurations, there is a one-to-one mapping between disks and storage groups. You can download the examples and edit the worksheets to plan the storage configuration for your server.

Chapter 6. Planning for security

Plan to protect the security of systems in the Tivoli Storage Manager solution with access and authentication controls, and consider encrypting data and password transmission.

Planning for administrator roles

Define the authority levels that you want to assign to administrators who have access to the Tivoli Storage Manager solution.

You can assign one of the following levels of authority to administrators:

System
Administrators who have system authority have the highest level of authority. Administrators with this level of authority can complete any task. They can manage all policy domains and storage pools, and grant authority to other administrators.

Policy Administrators who have policy authority can manage all of the tasks that are related to policy management. This privilege can be unrestricted, or can be restricted to specific policy domains.

Storage
Administrators who have Storage authority can allocate and control storage resources for the server.

Operator
Administrators who have Operator authority can control the immediate operation of the server and the availability of storage media such as tape libraries and drives.

The scenarios in Table 10 provide examples about why you might want to assign varying levels of authority so that administrators can perform tasks:

Table 10. Scenarios for administrator roles

Scenario	Type of administrator ID to set up
An administrator at a small company manages the server and is responsible for all server activities.	• System authority: 1 administrator ID
An administrator for multiple servers also manages the overall system. Several other administrators manage their own storage pools.	• System authority on all servers: 1 administrator ID for the overall system administrator • Storage authority for designated storage pools: 1 administrator ID for each of the other administrators
An administrator manages 2 servers. Another person helps with the administration tasks. Two assistants are responsible for helping to ensure that important systems are backed up. Each assistant is responsible for monitoring the scheduled backups on one of the Tivoli Storage Manager servers.	• System authority on both servers: 2 administrator IDs • Operator authority: 2 administrator IDs for the assistants with access to the server that each person is responsible for

Planning for secure communications

Plan for protecting communications among the solution components.

Determine the level of protection that is required for your data, based on regulations and business requirements under which your company operates.

If your business requires a high level of security for passwords and data transmission, plan on implementing secure communication with Transport Layer Security (TLS) or Secure Sockets Layer (SSL) protocols.

TLS and SSL provide secure communications between the server and client, but can affect system performance. TLS, a form of SSL, is required for all LDAP password communications. If you decide to use TLS or SSL, use the protocol only for sessions where it is necessary and add processor resources on the server to manage the increased requirements. You can also try other options, such as networking devices such as routers and switches that provide the TLS or SSL function.

You can use TLS and SSL to protect some or all of the different possible communication paths, for example:

- Operations Center: browser to hub; hub to spoke
- Client to server
- Server to server: node replication

Related tasks:

➡ Securing communications (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.admin/r_comms_securing.html)

Planning for storage of encrypted data

Determine whether your company requires stored data to be encrypted, and choose the option that best suits your needs.

If your company requires the data in storage pools to be encrypted, then you have the option of using Tivoli Storage Manager encryption, or an external device such as tape for encryption.

If you choose Tivoli Storage Manager to encrypt the data, extra computing resources are required at the client that might affect the performance for the backing up and restore processes.

Related information:

➡ Technote 1963635 (http://www.ibm.com/support/docview.wss?uid=swg21963635)

Planning firewall access

Determine the firewalls that are set and the ports that must be open for the Tivoli Storage Manager solution to work.

Table 11 describes the ports that are used by the Tivoli Storage Manager server, client, and Tivoli Storage Manager Operations Center.

Table 11. Ports that are used by the server, client, and Operations Center

Item	Default	Direction	Description
Base port (TCPPort)	1500	Outbound / Inbound	Each Tivoli Storage Manager server instance requires a unique TCP port. You can specify an alternative TCP port number. You can use the TCPADMINPORT option and ADMINONCLIENTPORT option to set TCP port values.
SSL Base Port (SSLTCPPort)	No default	Outbound / Inbound	This port is used only if SSL communication is enabled. A server can support both SSL and non-SSL communication if both TCPPort and SSLTCPPort are specified.
SMB	45	Inbound / Outbound	This port is used by configuration wizards that communicate by using native protocols with multiple hosts.
SSH	22	Inbound / Outbound	This port is used by configuration wizards that communicate by using native protocols with multiple hosts.
SMTP	25	Outbound	This port is used to send email alerts from the server.
NDMP	No default	Inbound / Outbound	The Tivoli Storage Manager server must be able to open an outbound NDMP control port connection to the NAS device. The outbound control port is the Low-Level Address in the data mover definition for the NAS device. During an NDMP filer-to-server restore, the server must be able to open an outbound NDMP data connection to the NAS device. The data connection port that is used during a restore can be configured on the NAS device. During NDMP filer-to-server backups, the NAS device must be able to open outbound data connections to the server and the server must be able to accept inbound NDMP data connections. You can use the server option NDMPPORTRANGE to restrict the set of ports available for use as NDMP data connections. You can configure a firewall for connections to these ports.

Table 11. Ports that are used by the server, client, and Operations Center (continued)

Item	Default	Direction	Description
Replication	No default	Outbound / Inbound	The port and protocol for the outbound port for replication are set by the define server command that is used to set up replication. The inbound ports for replication are the TCP ports and SSL ports that the source server names in the define server commands..
Client schedule port	Client port: 1501	Outbound	The client listens on the port that is named and communicates the port number to the server. The server contacts the client if server prompted scheduling is used. You can specify an alternative port number in the client options file.
Long running sessions	**KEEPALIVE** setting: YES	Outbound	When the **KEEPALIVE** option is enabled, keepalive packets are sent during client-server sessions to prevent the firewall software from closing long-running, inactive connections.
Operations Center	HTTPS: 11090	Inbound	These ports are used for the Operations Center web browser. You can specify an alternative port number.
Client management service port	Client port: 9028	Inbound	The client management service port must be accessible from the Tivoli Storage Manager Operations Center. Ensure there are no firewalls that can prevent connections. The client management service uses the TCP port of the client node's Tivoli Storage Manager server for authentication by using an administrative session.

Part 2. Tivoli Storage Manager multisite disk implementation

The multisite disk solution is configured at two sites and uses Tivoli Storage Manager data deduplication and replication.
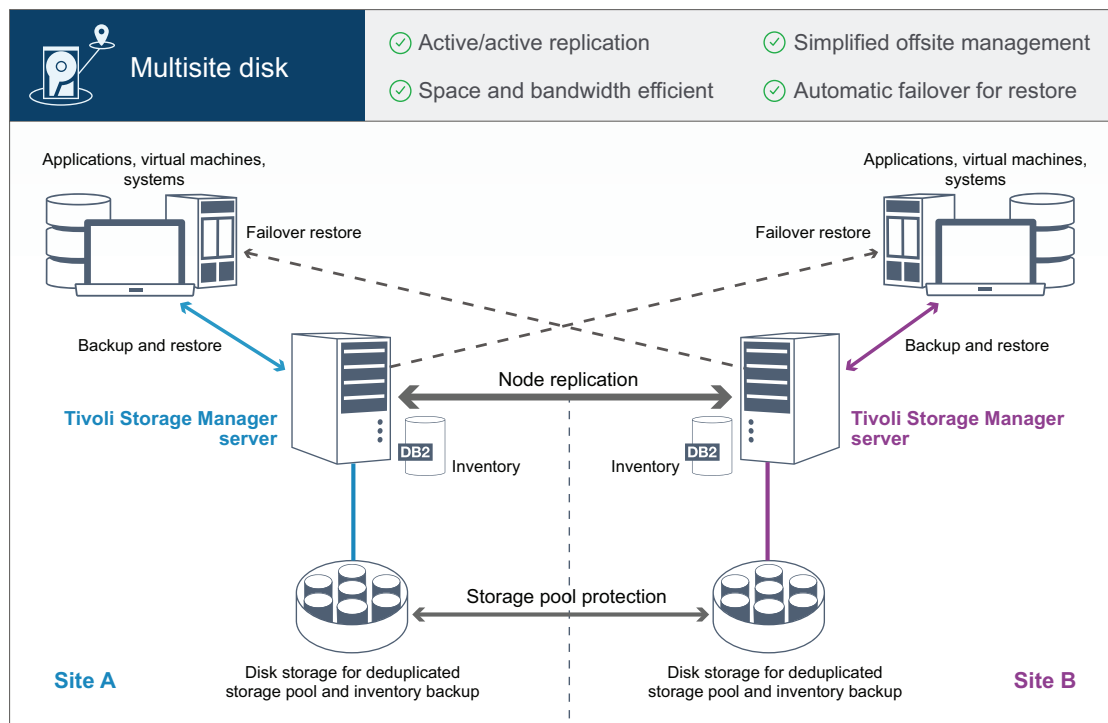
Implementation roadmap

The following steps are required to set up a Tivoli Storage Manager multisite disk environment.

1. Set up the system.
   a. Configure the storage hardware and set up storage arrays for your environment size.
   b. Install the server operating system.
   c. Configure multipath I/O.
   d. Create the user ID for the Tivoli Storage Manager server instance.
   e. Prepare file systems for Tivoli Storage Manager.
2. Install the Tivoli Storage Manager server and Operations Center.
3. Configure the Tivoli Storage Manager server and Operations Center.
   a. Complete the initial configuration of the server.
   b. Set server options.
   c. Configure Secure Sockets Layer for the server and client.
   d. Configure the Operations Center.
   e. Register your Tivoli Storage Manager license.
   f. Configure data deduplication.

      g. Define data retention rules for your business.

      h. Define server maintenance schedules.

      i. Define client schedules.

4. Install and configure clients.

      a. Register and assign clients to schedules.

      b. Install and verify the client management service.

      c. Configure the Operations Center to use the client management service.

5. Configure the second server.

      a. Confgure for SSL communication between the hub and spoke server.

      b. Add the second server as a spoke.

      c. Enable replication.

6. Complete the implementation.

Chapter 7. Setting up the system

To set up the system, you must first configure your disk storage hardware and the server system for Tivoli Storage Manager.

Configuring the storage hardware

To configure your storage hardware, review general guidance for disk systems and Tivoli Storage Manager.

Procedure

1. Provide a connection between the Tivoli Storage Manager server and the storage.
   - Use a switch or direct connection for Fibre Channel connections.
   - Consider the number of ports that will be connected and account for the amount of bandwidth that is needed.
   - Consider the number of ports on the Tivoli Storage Manager server and the number of host ports on the disk system that are connected.
2. Verify that device drivers and firmware for the server system, adapters, and operating system are current and at the recommended levels.
3. Configure storage arrays. Make sure that you have planned properly to ensure optimal performance. Refer to Chapter 5, "Planning for storage," on page 21 for more information.
4. The Tivoli Storage Manager server system needs access to disk volumes that are created. Ensure the following are completed:
   - If the system is connected to a Fibre Channel switch, zone the server to see the disks.
   - Map all of the volumes to tell the disk system that this specific server should be allowed to see each disk.

Installing the server operating system

Install the operating system on the server system and ensure that Tivoli Storage Manager server requirements are met. Adjust operating system settings as directed.

Installing on AIX systems

Complete the following steps to install AIX on the server system.

Procedure

1. Install AIX Version 7.1 TL3, SP4 or later according to the manufacturer instructions.
2. Configure your TCP/IP settings according to the operating system installation instructions.
3. Open the `/etc/hosts` file and complete the following actions:
   - Update the file to include the IP address and host name for the server. For example if your IP address is *192.0.2.7* and your server host name is *server.yourdomain.com*:

     ```
     192.0.2.7  server.yourdomain.com  server
     ```

- Verify that the file contains an entry for localhost with an address of 127.0.0.1. For example:

  ```
  127.0.0.1  localhost
  ```

4. Enable AIX I/O completion ports by issuing the following command:

   ```
   chdev -l iocp0 -P
   ```

5. Optional: Tivoli Storage Manager server performance can be affected by the Olson time zone definition. If performance is a factor in your environment, you can change your system time zone format from Olson to POSIX. Use the following command format to update the time zone setting:

   ```
   chtz=local_timezone,date/time,date/time
   ```

   For example, if you lived in Tucson, Arizona, where Mountain Standard Time is used, you would issue the following command to change to the POSIX format:

   ```
   chtz MST7MDT,M3.2.0/2:00:00,M11.1.0/2:00:00
   ```

   Also add an entry in the `.profile` of the instance user so that the following environment is set:

   ```
   export MALLOCOPTIONS=multiheap:16
   ```

6. For communications with the Tivoli Storage Manager server and Operations Center, make sure that the following ports are open on any firewalls that might exist:

   - Communications with the server: 1500
   - For secure communications with the Operations Center, open port 11090 on the hub server

   If you are not using the default port values, make sure that the ports you are using are open.

7. Verify that user limits, also known as *ulimits*, are set according to guidelines in Table 12. If ulimit values are not set correctly, you might experience server instability or a failure of the server to respond.

Table 12. User limits (ulimit) values

User limit type	Setting	Value	Command to query value
Maximum size of core files created	core	Unlimited	ulimit -Hc
Maximum size of a data segment for a process	data	Unlimited	ulimit -Hd
Maximum file size	fsize	Unlimited	ulimit -Hf
Maximum number of open files	nofile	65536	ulimit -Hn
Maximum amount of processor time in seconds	cpu	Unlimited	ulimit -Ht
Maximum number of user processes	nproc	16384	ulimit -Hu

If you need to modify any user limit values, follow the instructions in the documentation for your operating system.

Installing on Linux systems

Complete the following steps to install Linux x86_64 on the server system.

Before you begin

The operating system is installed on the internal hard disks. Configure the internal hard disks using a hardware RAID 1 array. For example, if you are configuring a small system, the two 300 GB internal disks are mirrored in RAID 1 so that a single 300 GB disk appears available to the operating system installer.

Procedure

1. Install Red Hat Enterprise Linux Version 7.1 or later, according to the manufacturer instructions. Obtain a bootable DVD that contains Red Hat Enterprise Linux Version 7.1 and start your system from this DVD. See the following guidance for installation options. If an item is not mentioned in the list, leave the default selection.

   a. After you start the DVD, choose **Install or upgrade an existing system** from the menu.

   b. On the Welcome screen, select **Test this media & install Red Hat Enterprise Linux 7.1**.

   c. Select your language and keyboard preferences.

   d. Select your location to set the correct time zone.

   e. Select **Software Selection** and then on the next screen, select **Server with GUI**.

   f. From the installation summary window, click **Installation Destination** and verify the following:

   - The local 300 GB disk is selected as the installation target.
   - Under Other Storage Options, `Automatically configure partitioning` is selected.

   Click **Done**.

   g. Click **Begin Installation**. After the installation starts, set the root password for your root user account.

   After the installation completes, restart the system and log in as the root user. Issue the `df` command to verify your basic partitioning. For example, on a test system, the initial partitioning appeared as follows:

   ```
   [root@tvapp02]# df –h
   Filesystem             Size  Used Avail Use% Mounted on
   /dev/mapper/rhel-root   50G  3.0G   48G   6% /
   devtmpfs                32G     0   32G   0% /dev
   tmpfs                   32G   92K   32G   1% /dev/shm
   tmpfs                   32G  8.8M   32G   1% /run
   tmpfs                   32G     0   32G   0% /sys/fs/cgroup
   /dev/mapper/rhel-home  220G   37M  220G   1% /home
   /dev/sda1              497M  124M  373M  25% /boot
   ```

2. Configure your TCP/IP settings according to the operating system installation instructions.

   For optimal throughput and reliability, consider bonding multiple network ports together. This can be accomplished by creating a Link Aggregation Control Protocol (LACP) network connection that aggregates several subordinate ports into a single logical connection. Configuration recommendations include using a bond mode of 802.3ad, `miimon` setting of 100, and a `xmit_hash_policy` setting of layer3+4.

For additional instructions on configuring bonded network connections with Red Hat Enterprise Linux Version 7, see https://access.redhat.com/ documentation/en-US/Red_Hat_Enterprise_Linux/7/html/ Networking_Guide/sec-Network_Bonding_Using_the_Command_Line_Interface.html#sec-Create_a_Channel_Bonding_Interface.

3. Open the /etc/hosts file and complete the following actions:
   - Update the file to include the IP address and host name for the server. For example:
     ```
     192.0.2.7  server.yourdomain.com  server
     ```
   - Verify that the file contains an entry for localhost with an address of 127.0.0.1. For example:
     ```
     127.0.0.1  localhost
     ```

4. Install components that are required for Tivoli Storage Manager server installation. Complete the following steps to create a Yellowdog Updater Modified (YUM) repository and install the prerequisite packages.
   a. Mount your Red Hat Enterprise Linux installation DVD to a system directory. For example, to mount it to the /mnt directory, issue the following command:
      ```
      mount -t iso9660 -o ro /dev/cdrom /mnt
      ```
   b. Verify that the DVD mounted by issuing the **mount** command. You should see output similar to the following:
      ```
      /dev/sr0 on /mnt type iso9660
      ```
   c. Change to the YUM repository directory by issuing the following command:
      ```
      cd /etc/yum/repos.d
      ```
      If the repos.d directory does not exist, create it.
   d. List directory contents:
      ```
      ls rhel-source.repo
      ```
   e. Rename the original repo file by issuing the **mv** command. For example:
      ```
      mv rhel-source.repo rhel-source.repo.orig
      ```
   f. Create a new repo file by using a text editor. For example, to use the vi editor, issue the following command:
      ```
      vi rhel71_dvd.repo
      ```
   g. Add the following lines to the new repo file. The **baseurl** parameter specifies your directory mount point:
      ```
      [rhel71_dvd]
      name=DVD Redhat Enterprise Linux 7.1
      baseurl=file:///mnt
      enabled=1
      gpgcheck=0
      ```
   h. Install the prerequisite package ksh.x86_64, which is required for Tivoli Storage Manager, by issuing the **yum** command. For example:
      ```
      yum install ksh.x86_64
      ```
      Note: You do not need to install the compat-libstdc++-33-3.2.3-69.el6.i686 and libstdc++.i686 libraries for Red Hat Enterprise Linux Version 7.1.

5. When the software installation is complete, you can restore the original YUM repository values by completing the following steps:

a. Unmount the Red Hat Enterprise Linux installation DVD by issuing the following command:

```
unmount /mnt
```

b. Change to the YUM repository directory by issuing the following command:

```
cd /etc/yum/repos.d
```

c. Rename the repo file that you created:

```
mv rhel71_dvd.repo rhel71_dvd.repo.orig
```

d. Rename the original file to the original name:

```
mv rhel-source.repo.orig rhel-source.repo
```

6. Determine if kernel parameter changes are required. Complete the following steps:

a. Use the **sysctl -a** command to list the parameter values.

b. Analyze the results by using the guidelines in Table 13 to determine if any changes are required.

c. If changes are required, set the parameters in the /etc/sysctl.conf file. The file changes are applied when the system is started.

Note: Kernel parameter values that are listed in Table 13 include commas for ease of reading. Do not include commas for any value that you update in the /etc/sysctl.conf file.

Note: On Linux, DB2 might automatically increase interprocess communication (IPC) kernel parameter values to the preferred settings. If DB2 updates the values that you set, you do not need to change them back to the values that are listed in the following table.

Table 13. Linux kernel parameter optimum settings

Parameter	Description	Preferred value
kernel.shmmni	The maximum number of segments.	256 x *size of RAM in GB* Values for each system size: • Small: 16,384 • Medium: 32,768 • Large: 49,152
kernel.shmmax	The maximum size of a shared memory segment (bytes). This parameter must be set before automatically starting the Tivoli Storage Manager server on system startup.	*Size of RAM in bytes* Values for each system size: • Small: 68,719,476,736 • Medium: 137,438,953,472 • Large: 206,158,430,208
kernel.shmall	The maximum allocation of shared memory pages (pages)	2 x *size of RAM in bytes* (setting is in 4 KB pages) Value used for all sized systems: 4,294,967,296 Changes to the factory settings for this parameter are not required.

Table 13. Linux kernel parameter optimum settings (continued)

Parameter	Description	Preferred value

kernel.sem
 	
You must specify four values for the kernel.sem parameter. When you update this parameter, include all values on one line in the following order:
`kernel.sem = SEMMSL SEMMNS SEMOPM SEMMNI`
 	
For example, to update the parameter for a medium system, enter the following on one line in the /etc/sysctl.conf file:
`kernel.sem = 250 256000 32 32768`
	(SEMMSL)
	
	The maximum semaphores per array
		250

	(SEMMNS) The maximum semaphores per system	256,000
	(SEMOPM) The maximum operations per semaphore call	32
	(SEMMNI) The maximum number of arrays	256 x *size of RAM in GB* Values for each system size: • Small: 16,384 • Medium: 32,768 • Large: 49,152
kernel.msgmni	The maximum number of system-wide message queues	1024 x *size of RAM in GB* Values for each system size: • Small: 65,536 • Medium: 131,072 • Large: 196,608
kernel.msgmax	The maximum size of messages (bytes)	65,536
kernel.msgmnb	The default maximum size of queue (bytes)	65,536
kernel.randomize_va_space	The **kernel.randomize_va_space** parameter configures the use of memory ASLR for the kernel. Disable ASLR because it can cause errors for the DB2 software. To learn more details about the Linux ASLR and DB2, see Technote 1365583 at: http://www.ibm.com/support/docview.wss?uid=swg21365583.	0
vm.swappiness	The **vm.swappiness** parameter defines whether the kernel can swap application memory out of physical random access memory (RAM). For more information about kernel parameters, see the DB2 information center.	0
vm.overcommit_memory	The **vm.overcommit_memory** parameter influences how much virtual memory the kernel can permit to be allocated. For more information about kernel parameters, see the DB2 information center.	0

7. Open firewall ports that Tivoli Storage Manager will use. Complete the following steps:

   a. Determine the zone that is used by the network interface. The zone is public, by default.

      Issue the following command:

```
#  firewall-cmd --get-active-zones
public
  interfaces: ens4f0
```

    b. To use the default port address for communications with the Tivoli Storage
       Manager server, open TCP/IP port 1500 in the Linux firewall.

       Issue the following command:

```
firewall-cmd --zone=public --add-port=1500/tcp --permanent
```

       If you want to use a value other than the default, you can specify a number
       in the range 1024 - 32767. If you open a port other than the default, you
       will need to specify that port when you run the configuration script.

    c. If you plan to use this system as a Tivoli Storage Manager Operations
       Center hub, open port 11090, which is the default port for secure (https)
       communications.

       Issue the following command:

```
firewall-cmd --zone=public --add-port=11090/tcp --permanent
```

    d. Reload the firewall definitions for the changes to take effect.

       Issue the following command:

```
firewall-cmd --reload
```

8. Verify that user limits, also known as *ulimits*, are set according to guidelines in
   Table 14. If ulimit values are not set correctly, you might experience server
   instability or a failure of the server to respond.

Table 14. User limits (ulimit) values

User limit type	Setting	Value	Command to query value
Maximum size of core files created	core	Unlimited	ulimit -Hc
Maximum size of a data segment for a process	data	Unlimited	ulimit -Hd
Maximum file size	fsize	Unlimited	ulimit -Hf
Maximum number of open files	nofile	65536	ulimit -Hn
Maximum amount of processor time in seconds	cpu	Unlimited	ulimit -Ht
Maximum number of user processes	nproc	16384	ulimit -Hu

   If you need to modify any user limit values, follow the instructions in the
   documentation for your operating system.

Installing on Windows systems

Install Microsoft Windows Server 2012 Standard Edition on the server system and
prepare the system for installation and configuration of the Tivoli Storage Manager
server.

Procedure

1. Install Microsoft Windows Server 2012 R2 Standard Edition, according to the
   manufacturer instructions.
2. Change the Windows account control policies by completing the following
   steps.
    a. Start the Local Security Policy editor by running secpol.msc.

b. Go to **Local Policies** > **Security Options** and ensure that the following User Account Control policies are disabled:

- Admin Approval Mode for the Built-in Administrator account
- Run all administrators in Admin Approval Mode

3. Configure your TCP/IP settings according to installation instructions for the operating system.

4. Apply Windows updates and enable optional features by completing the following steps:

a. Apply the latest Windows 2012 updates.

b. Install and enable the Windows 2012 feature Microsoft .NET Framework 3.5 from the Windows Server Manager.

c. Update HBA device drivers as needed.

d. Install the multipath I/O driver that is appropriate for the disk system that you are using.

5. Open the default TCP/IP port, 1500, for communications with the Tivoli Storage Manager server. For example, issue the following command:

```
netsh advfirewall firewall add rule name="TSM server port 1500"
dir=in action=allow protocol=TCP localport=1500
```

6. On the Operations Center hub server, open the default port for secure (https) communications with the Operations Center. The port number is 11090. For example, issue the following command

```
netsh advfirewall firewall add rule name="TSM Operations Center port 11090"
dir=in action=allow protocol=TCP localport=11090
```

Configuring multipath I/O

Complete the following steps to enable and configure multipathing for disk storage. Use the documentation that is provided with your hardware for detailed instructions.

AIX systems

Procedure

1. Determine the Fibre Channel port address that you must use for the host definition on the disk subsystem. Issue the **lscfg** command for every port.

- On small and medium systems, issue the following commands:

```
lscfg -vps -l fcs0 | grep "Network Address"
lscfg -vps -l fcs1 | grep "Network Address"
```

- On large systems, issue the following commands:

```
lscfg -vps -l fcs0 | grep "Network Address"
lscfg -vps -l fcs1 | grep "Network Address"
lscfg -vps -l fcs2 | grep "Network Address"
lscfg -vps -l fcs3 | grep "Network Address"
```

2. Ensure that the following AIX file sets are installed:

- devices.common.IBM.mpio.rte
- devices.fcp.disk.array.rte
- devices.fcp.disk.rte

3. Issue the **cfgmgr** command to have AIX rescan the hardware and discover available disks. For example:

```
cfgmgr
```

4. To list the available disks, issue the following command:

```
lsdev -Ccdisk
```

You should see output similar to the following:

```
hdisk0  Available 00-00-00 SAS Disk Drive
hdisk1  Available 00-00-00 SAS Disk Drive
hdisk2  Available 01-00-00 SAS Disk Drive
hdisk3  Available 01-00-00 SAS Disk Drive
hdisk4  Available 06-01-02 MPIO IBM 2076 FC Disk
hdisk5  Available 07-01-02 MPIO IBM 2076 FC Disk
...
```

5. Use the output from the **lsdev** command to identify and list device IDs for each disk device.

   For example, hdisk4. Save the list of device IDs to use when you create file systems for the Tivoli Storage Manager server.

6. Correlate the SCSI device IDs to specific disk LUNs from the disk system by listing detailed information about all physical volumes in the system. Issue the following command:

```
lspv -u
```

For example, on an IBM Storwize system, you should see output similar to the following for each device:

```
hdisk4  00f8cf083fd97327  None  active
 332136005076300810105780000000000003004214503IBMfcp
```

In the example, *6005076300810105780000000000030* is the UID for the volume, as reported by the Storwize management interface.

To verify disk size in MB and compare with what is listed for the system, issue the following command:

```
bootinfo -s hdisk4
```

Linux systems

Procedure

1. Edit the /etc/multipath.conf file to enable multipathing for Linux hosts. For example, the following parameters were set in multipath.conf for testing on an IBM Storwize system:

```
defaults {
        user_friendly_names no
}

devices {
        device {
                vendor "IBM "
                product "2145"
                path_grouping_policy group_by_prio
                user_friendly_names no
        }
}
```

If the multipath.conf file does not exist, you can create it by issuing the following command:

```
mpathconf --enable
```

2. Set the multipath option to start when the system is started. Issue the following commands:

```
systemctl enable multipathd.service
systemctl start multipathd.service
```

3. To verify that disks are visible to the operating system and are managed by multipath, issue the following command:

```
multipath -l
```

4. Ensure that each device is listed and that it has as many paths as you expect. You can use size and device ID information to identify which disks are listed.

For example, on an IBM Storwize system, the following output shows that a 2 TB disk has two path groups and four active paths. The 2 TB size confirms that the disk corresponds to a pool file system. Use part of the long device ID number (12, in this example) to search for the volume on the disk-system management interface.

```
[root@tapsrv01 code]# multipath -l
36005076802810c509800000000000012 dm-43 IBM,2145
 size=2.0T features='1 queue_if_no_path' hwhandler='0' wp=rw
|-+- policy='round-robin 0' prio=0 status=active
| |- 2:0:1:18 sdcw 70:64   active undef running
| `- 4:0:0:18 sdgb 131:112 active undef running
`-+- policy='round-robin 0' prio=0 status=enabled
  |- 1:0:1:18 sdat 66:208  active undef running
  `- 3:0:0:18 sddy 128:0   active undef running
```

   a. If needed, correct disk LUN host assignments and force a bus rescan. For example:

```
echo "- - -" > /sys/class/scsi_host/host0/scan
echo "- - -" > /sys/class/scsi_host/host1/scan
echo "- - -" > /sys/class/scsi_host/host2/scan
```

   You can also restart the system to rescan disk LUN host assignments.

   b. Confirm that disks are now available for multipath I/O by reissuing the **multipath -l** command.

5. Use the multipath output to identify and list device IDs for each disk device.

For example, the device ID for your 2 TB disk is 36005076802810c509800000000000012.

Save the list of device IDs to use in the next step.

Windows systems

Procedure

1. Ensure that the Multipath I/O feature is installed. If needed, install additional vendor-specific multipath drivers.

2. To verify that disks are visible to the operating system and are managed by multipath I/O, issue the following command:

```
c:\program files\IBM\SDDDSM\datapath.exe query device
```

3. Review the multipath output and ensure that each device is listed and that it has as many paths as you expect. You can use size and device serial information to identify which disks are listed.

For example, by using part of the long device serial number (34, in this example) you can search for the volume on the disk-system management interface. The 2 TB size confirms that the disk corresponds to a storage pool file system.

```
DEV#:   4  DEVICE NAME: Disk5 Part0   TYPE: 2145    POLICY: OPTIMIZED
SERIAL: 6005076300810105780000000000000034          LUN SIZE: 2.0TB
====================================================================
Path#          Adapter/Hard Disk      State   Mode   Select  Errors
    0   Scsi Port2 Bus0/Disk5 Part0    OPEN    NORMAL      0       0
```

```
        1   Scsi Port2 Bus0/Disk5 Part0      OPEN    NORMAL   27176       0
        2   Scsi Port3 Bus0/Disk5 Part0      OPEN    NORMAL   28494       0
        3   Scsi Port3 Bus0/Disk5 Part0      OPEN    NORMAL       0       0
```

4. Create a list of disk device IDs by using the serial numbers that are returned from the multipath output in Step 3.

   For example, the device ID for your 2 TB disk is
   600507630081010578000000000000034

   Save the list of device IDs to use in the next step.

5. After the new disks are added, you might need to bring them on-line and clear the read-only attribute. Run `diskpart.exe` with the following commands. Repeat for each of the disks:

```
diskpart
   select Disk 1
   online disk
   attribute disk clear readonly
   select Disk 2
   online disk
   attribute disk clear readonly
   < ... >
   select Disk 49
   online disk
   attribute disk clear readonly
   exit
```

Creating the user ID for the Tivoli Storage Manager server instance

Create the user ID that will own the Tivoli Storage Manager server instance. You will specify this user ID when you create the server instance during initial configuration of the server.

About this task

You can only specify lowercase letters (a-z), numerals (0-9), and the underscore character (_) for the user ID. The user ID and group name must comply with the following rules:

- The length must be 8 characters or less.
- The user ID and group name cannot start with *ibm, sql, sys,* or a numeral.
- The user ID and group name cannot be *user, admin, guest, public, local,* or any SQL reserved word.

Procedure

1. Use operating system commands to create a user ID.

   - AIX   Linux Create a group and user ID in the home directory of the user that will own the Tivoli Storage Manager server instance.

     For example, to create the user ID tsminst1 in group tsmsrvrs with a password of tsminst1, issue the following commands from an administrative user ID:

      AIX

```
mkgroup id=1001 tsmsrvrs
mkuser id=1002 pgrp=tsmsrvrs home=/home/tsminst1 tsminst1
passwd tsminst1
```

      Linux

```
groupadd tsmsrvrs
useradd -d /home/tsminst1 -m -g tsmsrvrs -s /bin/bash tsminst1
passwd tsminst1
```

Log off, and then log in to your system. Change to the user account that you just created. Use an interactive login program, such as telnet, so that you are prompted for the password and can change it if necessary.

* Windows Create a user ID and then add the new ID to the Administrators groups. For example, to create the user ID tsminst1, issue the following command:

```
net user tsminst1 * /add
```

After you create and verify a password for the new user, add the user ID to the Administrators group by issuing the following commands:

```
net localgroup Administrators tsminst1 /add
net localgroup DB2ADMNS tsminst1 /add
```

2. Log off the new user ID.

Preparing file systems for Tivoli Storage Manager

You must complete file system configuration for your disk storage to be used with Tivoli Storage Manager.

AIX systems

You must create volume groups, logical volumes, and file systems for the Tivoli Storage Manager server by using the AIX Logical Volume Manager.

Procedure

1. Increase the queue depth and maximum transfer size for all of the available *hdiskX* disks that were listed in the previous step. Issue the following commands for each disk:

```
chdev -l hdisk4 -a max_transfer=0x100000
chdev -l hdisk4 -a queue_depth=32
chdev -l hdisk4 -a reserve_policy=no_reserve
chdev -l hdisk4 -a algorithm=round_robin
```

Do not run these commands for operating system internal disks, for example, *hdisk0*.

2. Create volume groups for the Tivoli Storage Manager database, active log, archive log, database backup, and storage pool. Issue the **mkvg** command, specifying the device IDs for corresponding disks that you previously identified.

For example, if the device names *hdisk4*, *hdisk5*, and *hdisk6* correspond to database disks, include them in the database volume group and so on.

Note: The following commands are based on the medium system configuration. The syntax issued for small and large systems will require minor adjustments.

```
mkvg -S -y tsmdb hdisk2 hdisk3 hdisk4
mkvg -S -y tsmactlog hdisk5
mkvg -S -y tsmarchlog hdisk6
mkvg -S -y tsmdbback hdisk7 hdisk8 hdisk9 hdisk10
mkvg -S -y tsmstgpool hdisk11 hdisk12 hdisk13 hdisk14 ... hdisk49
```

3. Determine the physical volume names and the number of free physical partitions to use when you create logical volumes. Issue the **lsvg** for each volume group that you created in the previous step.

   For example:

   ```
   lsvg -p tsmdb
   ```

   The output is similar to the following. The *FREE PPS* column represents the free physical partitions:

   ```
   tsmdb:
   PV_NAME  PV STATE   TOTAL PPs  FREE PPs  FREE DISTRIBUTION
   hdisk4   active     1631       1631      327..326..326..326..326
   hdisk5   active     1631       1631      327..326..326..326..326
   hdisk6   active     1631       1631      327..326..326..326..326
   ```

4. Create logical volumes in each volume group by using the **mklv** command. The volume size, volume group, and device names vary, depending on the size of your system and variations in your disk configuration.

   For example, to create the volumes for the Tivoli Storage Manager database on a medium system, issue the following commands:

   ```
   mklv -y tsmdb00 -t jfs2 -u 1 -x 1631 tsmdb 1631 hdisk2
   mklv -y tsmdb01 -t jfs2 -u 1 -x 1631 tsmdb 1631 hdisk3
   mklv -y tsmdb02 -t jfs2 -u 1 -x 1631 tsmdb 1631 hdisk4
   ```

5. Format file systems in each logical volume by using the **crfs** command.

   For example, to format file systems for the database on a medium system, issue the following commands:

   ```
   crfs -v jfs2 -d tsmdb00 -p rw -a logname=INLINE -a options=rbrw
    -a agblksize=4096 -m /tsminst1/TSMdbspace00 -A yes
   crfs -v jfs2 -d tsmdb01 -p rw -a logname=INLINE -a options=rbrw
    -a agblksize=4096 -m /tsminst1/TSMdbspace01 -A yes
   crfs -v jfs2 -d tsmdb02 -p rw -a logname=INLINE -a options=rbrw
    -a agblksize=4096 -m /tsminst1/TSMdbspace02 -A yes
   ```

6. Mount all of the newly created file systems by issuing the following command:

   ```
   mount -a
   ```

7. List all file systems by issuing the **df** command. Verify that file systems are mounted at the correct LUN and correct mount point. Also, verify the available space. The amount of used space should be around one percent.

   For example:

   ```
   tapsrv07> df -g /tsminst1/*
   Filesystem      GB blocks  Free     %Used   Iused   %Iused   Mounted on
   /dev/tsmact00   195.12     194.59   1%      4       1%       /tsminst1/TSMalog
   ```

8. Verify that the user ID you created in "Creating the user ID for the Tivoli Storage Manager server instance" on page 39 has read and write access to the directories for Tivoli Storage Manager.

Linux systems

You must format ext4 or xfs file systems on each of the disk LUNs that the Tivoli Storage Manager server will use.

Procedure

1. Using the list of device IDs that you generated in the previous step, issue the **mkfs** command to create and format a file system for each storage LUN device. Specify the device ID in the command. For example: For the database, format ext4 file systems:

   ```
   mkfs -t ext4 -T largefile -m 2 /dev/mapper/36005076802810c5098000000000000012
   ```

For storage pool LUNs, format xfs file systems:

```
mkfs -t xfs /dev/mapper/36005076300810105780000000000002c3
```

You might issue the **mkfs** command as many as 50 times, depending on how many different devices you have.

2. Create mount point directories for Tivoli Storage Manager file systems.

   Issue the **mkdir** command for each Tivoli Storage Manager directory that you must create. Use the directory values that you recorded in the planning worksheets. For example, to create the server instance directory by using the default value, issue the following command:

```
mkdir /tsminst1
```

   Repeat the **mkdir** command for each file system.

3. Add an entry in the /etc/fstab file for each file system so that file systems are mounted automatically when the server is started.

   For example:

```
/dev/mapper/36005076802810c5098000000000000012 /tsminst1/TSMdbspace00 ext4 defaults 0 0
```

4. Mount the file systems that you added to the /etc/fstab file by issuing the **mount -a** command.

5. List all file systems by issuing the **df** command. Verify that file systems are mounted at the correct LUN and correct mount point. Also, verify the available space. The amount of used space should be around one percent.

   For example, on an IBM Storwize system, the following output is returned:

```
[root@tapsrv04 ~]# df -h /tsminst1/*
Filesystem                                      Size  Used Avail Use% Mounted on
/dev/mapper/36005076300810105780000000000003 134G  188M 132G  1%   /tsminst1/TSMalog
```

6. Verify that the user ID you created in "Creating the user ID for the Tivoli Storage Manager server instance" on page 39 has read and write access to the directories for Tivoli Storage Manager.

Windows systems

You must format NTFS file systems on each of the disk LUNs that the Tivoli Storage Manager server will use.

Procedure

1. Create mount point directories for Tivoli Storage Manager file systems.

   Issue the **md** command for each Tivoli Storage Manager directory that you must create. Use the directory values that you recorded in the planning worksheets. For example, to create the server instance directory by using the default value, issue the following command:

```
md c:\tsminst1
```

   Repeat the **md** command for each file system.

2. Create a volume for every disk LUN that is mapped to a directory under the server instance directory by using the Windows volume manager.

   Go to **Server Manager** > **File and Storage Services** and complete the following steps for each disk that corresponds to the LUN mapping that was created in the previous step:

   a. Bring the disk online.

   b. Initialize the disk to the GPT basic type, which is the default.

   c. Create a simple volume that occupies all of the space on the disk. Format the file system by using NTFS, and assign a label that matches the purpose

of the volume, such as `TSMfile00`. Do not assign the new volume to a drive letter. Instead, map the volume to a directory under the instance directory, such as `C:\tsminst1\TSMfile00`.

> **Tip:** Determine the volume label and directory mapping labels based on the size of the disk that is reported.

3. Verify that file systems are mounted at the correct LUN and correct mount point. List all file systems by issuing the **mountvol** command and then review the output. For example:

```
\\?\Volume{8ffb9678-3216-474c-a021-20e420816a92}\
   C:\tsminst1\TSMdbspace00\
```

4. After the disk configuration is complete, restart the system.

What to do next

You can confirm the amount of free space for each volume by using Windows Explorer.

Chapter 8. Installing the Tivoli Storage Manager server and Operations Center

Use the IBM Installation Manager graphical wizard to install Tivoli Storage Manager components.

Installing on AIX and Linux systems

Install the Tivoli Storage Manager server and the Operations Center on the first server system.

Before you begin

Verify that the operating system is set to the language that you require. By default, the language of the operating system is the language of the installation wizard.

Procedure

1. `AIX` Verify that the required RPM files are installed on your system.

   See "Installing prerequisite RPM files for the graphical wizard" on page 46 for details.

2. Before you download the installation package, verify that you have enough space to store the installation files when they are extracted from the product package. Refer to the download document for space requirements: http://www.ibm.com/support/docview.wss?uid=swg24035122

3. Go to Passport Advantage® at: http://www.ibm.com/software/passportadvantage/. Download the package file to an empty directory of your choice.

4. Ensure that executable permission is set for the package. If necessary, change the file permissions by issuing the following command:

   `chmod a+x package_name.bin`

5. Extract the package by issuing the following command:

   `./package_name.bin`

   where *package_name* is the name of the downloaded file.

6. `AIX` Ensure that the following command is enabled so that the Tivoli Storage Manager wizards work properly:

   `lsuser`

   By default, the command is enabled.

7. Change to the directory where you placed the executable file.

8. Start the installation wizard by issuing the following command:

   `./install.sh`

   When you select the packages to install, choose both the server and Operations Center.

What to do next

- If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM Installation Manager logs directory.

You can view installation log files by clicking **File > View Log** from the Installation Manager tool. To collect these log files, click **Help > Export Data for Problem Analysis** from the Installation Manager tool.

- After you install Tivoli Storage Manager, and before you customize it for your use, go to the following website: http://www.ibm.com/support/entry/portal/ Overview/Software/Tivoli/Tivoli_Storage_Manager. Click **Support and downloads** and apply any applicable fixes.

Related tasks:

➦ Other methods for installing Tivoli Storage Manager components (AIX) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.install/ t_srv_inst-aix.html)

➦ Other methods for installing Tivoli Storage Manager components (Linux) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.install/ t_srv_inst-linux.html)

Installing prerequisite RPM files for the graphical wizard

AIX

RPM files are required for the IBM Installation Manager graphical wizard.

Procedure

1. Verify that the following files are installed on your system. If the files are not installed, go to Step 2.

   ```
   atk-1.12.3-2.aix5.2.ppc.rpm               libpng-1.2.32-2.aix5.2.ppc.rpm
   cairo-1.8.8-1.aix5.2.ppc.rpm              libtiff-3.8.2-1.aix5.2.ppc.rpm
   expat-2.0.1-1.aix5.2.ppc.rpm              pango-1.14.5-4.aix5.2.ppc.rpm
   fontconfig-2.4.2-1.aix5.2.ppc.rpm         pixman-0.12.0-3.aix5.2.ppc.rpm
   freetype2-2.3.9-1.aix5.2.ppc.rpm          xcursor-1.1.7-3.aix5.2.ppc.rpm
   gettext-0.10.40-6.aix5.1.ppc.rpm          xft-2.1.6-5.aix5.1.ppc.rpm
   glib2-2.12.4-2.aix5.2.ppc.rpm             xrender-0.9.1-3.aix5.2.ppc.rpm
   gtk2-2.10.6-4.aix5.2.ppc.rpm              zlib-1.2.3-3.aix5.1.ppc.rpm
   libjpeg-6b-6.aix5.1.ppc.rpm
   ```

2. Ensure that there is at least 150 MB of free space in the /opt file system.

3. From the directory where the installation package file is extracted, go to the gtk directory.

4. Download the RPM files to the current working directory from the IBM AIX Toolbox for Linux Applications website (http://www.ibm.com/systems/ power/software/aix/linux/) by issuing the following command:

   ```
   download-prerequisites.sh
   ```

5. From the directory that contains the RPM files that you downloaded, install them by issuing the following command:

   ```
   rpm -Uvh *.rpm
   ```

Installing on Windows systems

Install the Tivoli Storage Manager server and the Operations Center on the first server system.

Before you begin

Make sure that the following prerequisites are met:
- Verify that the operating system is set to the language that you require. By default, the language of the operating system is the language of the installation wizard.
- Ensure that the user ID that you plan to use during the installation is a user with local Administrator authority.

Procedure

1. Before you download the installation package, verify that you have enough space to store the installation files when they are extracted from the product package. Refer to the download document for space requirements: http://www.ibm.com/support/docview.wss?uid=swg24035122
2. Go to Passport Advantage at: http://www.ibm.com/software/passportadvantage/. Download the package file to an empty directory of your choice.
3. Change to the directory where you placed the executable file.
4. Double-click the executable file to extract to the current directory.
5. In the directory where the installation files were extracted, start the installation wizard by double-clicking the `install.bat` file. When you select the packages to install, choose both the server and Operations Center.

What to do next

- If errors occur during the installation process, the errors are recorded in log files that are stored in the IBM Installation Manager logs directory.

  You can view installation log files by clicking **File > View Log** from the Installation Manager tool. To collect these log files, click **Help > Export Data for Problem Analysis** from the Installation Manager tool.
- After you install Tivoli Storage Manager, and before you customize it for your use, go to the following website: http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Manager. Click **Support and downloads** and apply any applicable fixes.

Related tasks:

➡ Other methods for installing Tivoli Storage Manager components (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.install/t_srv_inst-windows.html)

Chapter 9. Configuring the Tivoli Storage Manager server and the Operations Center

After you install Tivoli Storage Manager, complete the configuration for the server and the Operations Center.

Configuring the server instance

Use the Tivoli Storage Manager server instance configuration wizard to complete the initial configuration of the server.

Before you begin

Ensure that the following requirements are met:

AIX   Linux

- The system where you installed Tivoli Storage Manager must have the X Window System client. You must also be running an X Window System server on your desktop.
- The system must have the Secure Shell (SSH) protocol enabled. Ensure that the port is set to the default value, 22, and that the port is not blocked by a firewall. You must enable password authentication in the sshd_config file in the /etc/ssh/directory. Also, ensure that the SSH daemon service has access rights to connect to the system by using the *localhost* value.
- You must be able to log in to Tivoli Storage Manager with the user ID that you created for the server instance, by using the SSH protocol. When you use the wizard, you must provide this user ID and password to access that system.
- If you make changes, restart the server before you proceed with the configuration wizard.

Windows   Verify that the remote registry service is started by completing the following steps:

1. Click **Start > Administrative Tools > Services**. In the Services window, select **Remote Registry**. If it is not started, click **Start**.
2. Ensure port 137, 139 and 445 are not blocked by a firewall:
   a. Click **Start > Control Panel > Windows Firewall**.
   b. Select **Advanced Settings**.
   c. Select **Inbound Rules** in the left pane.
   d. Select **New Rule** in the right pane.
   e. Create a port rule for TCP ports 137, 139 and 445 to allow connections for domain and private networks.
3. Configure the user account control by accessing the local security policy options and completing the following steps.
   a. Click **Start** > **Administrative Tools** > **Local Security Policy**. Expand **Local Policies** > **Security Options**.
   b. If not already enabled, enable the built-in administrator account by selecting **Accounts: Administrator account status > Enable > OK**.

c. If not already disabled, disable user account control for all Windows administrators by selecting **User Account Control: Run all administrators in Admin Approval Mode > Disable > OK**.

d. If not already disabled, disable the User Account Control for the built-in Administrator account by selecting **User Account Control: Admin Approval Mode for the Built-in Administrator Account >Disable > OK**.

4. If you make changes, restart the server before you proceed with the configuration wizard.

About this task

The wizard can be stopped and restarted, but the server is not operational until the entire configuration process is complete.

Procedure

1. Start the local version of the wizard.

   - AIX   Linux   Open the `dsmicfgx` program in the `/opt/tivoli/tsm/server/bin` directory. This wizard can be only run as a root user.

   - Windows   Click **Start > All Programs > Tivoli Storage Manager > Configuration Wizard**.

2. Follow the instructions to complete the configuration. Use the information that you recorded in Chapter 4, "Planning worksheets," on page 11 during Tivoli Storage Manager system set up to specify directories and options in the wizard.

   AIX   Linux   On the Server Information window, set the server to start automatically by using the instance user ID when the machine boots.

   Windows   By using the configuration wizard, the server is set to start automatically when rebooted.

Installing the Tivoli Storage Manager backup-archive client

As a best practice, install the Tivoli Storage Manager backup-archive client on the server system so that the administrative command-line client and scheduler are available.

Procedure

To install the Tivoli Storage Manager backup-archive client, follow the installation instructions for your operating system.

- Installing the UNIX and Linux backup-archive clients (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/client/c_inst_baunix.html)
- Installing the Windows backup-archive client (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/client/t_inst_winclient.html)

Setting options for the server

Review the server options file that is installed with the Tivoli Storage Manager server, to verify that the correct values are set for your system.

Procedure

1. Go to the server instance directory and open the `dsmserv.opt` file.
2. Review the following table values and verify your server option settings, based on system size.

Server option	Small system value	Medium system value	Large system value
`ACTIVELOGDIRECTORY`	Directory path that was specified during configuration	Directory path that was specified during configuration	Directory path that was specified during configuration
`ACTIVELOGSIZE`	131072	131072	262144
`ARCHLOGCOMPRESS`	Yes	No	No
`ARCHLOGDIRECTORY`	Directory path that was specified during configuration	Directory path that was specified during configuration	Directory path that was specified during configuration
`COMMMETHOD`	TCPIP	TCPIP	TCPIP
`COMMTIMEOUT`	3600	3600	3600
`DEDUPREQUIRESBACKUP`	No	No	No
`DEVCONFIG`	`devconf.dat`	`devconf.dat`	`devconf.dat`
`EXPINTERVAL`	0	0	0
`IDLETIMEOUT`	60	60	60
`MAXSESSIONS`	250	500	1000
`NUMOPENVOLSALLOWED`	20	20	20
`TCPADMINPORT`	1500	1500	1500
`TCPPORT`	1500	1500	1500
`VOLUMEHISTORY`	`volhist.dat`	`volhist.dat`	`volhist.dat`

Update server option settings if necessary, to match the values in the table. If you have to make changes, close the `dsmserv.opt` file and use the **SETOPT** command from the administrative command-line interface to set the options.

For example, to update the `IDLETIMEOUT` option to 60, issue the following command:

```
setopt idletimeout 60
```

3. If you have to configure secure communications for the server, clients, and the Operations Center, verify the options in the following table.

Server option	All sized systems
`SSLDISABLELEGACYTLS`	YES
`SSLFIPSMODE`	NO
`SSLTCPPORT`	Specify the SSL port number. The server TCP/IP communication driver waits for requests on this port for SSL-enabled sessions from the client.

Server option	All sized systems
SSLTCPADMINPORT	Specify the port address on which the server waits for requests for SSL-enabled sessions from the command-line administrative client.
SSLTLS12	YES

If any of the option values must be updated, edit the `dsmserv.opt` file by using the following guidelines:

- Remove the asterisk at the beginning of a line to enable an option.
- Enter only one option and the specified value for the option per line.
- If there are multiple entries for an option in the file, the server uses the last entry.

Save your changes and close the file. If you edit the `dsmserv.opt` file directly, you must restart the server for the changes to take effect.

Related reference:

⬆ Server options reference (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_opt_server.html)

⬆ SETOPT (Set a server option for dynamic update) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_setopt.html)

Configuring secure communications with Transport Layer Security

If your environment requires secure communications, you can set up Secure Sockets Layer (SSL) and Transport Layer Security (TLS) on the Tivoli Storage Manager server and backup-archive client to encrypt data. An SSL certificate is used to verify communication requests between the server and client.

About this task

As shown in the following figure, you can configure SSL/TLS communications between the Tivoli Storage Manager server and backup-archive client by setting options in the server and client options files, and then transferring the self-signed certificate that is generated on the server to the client.
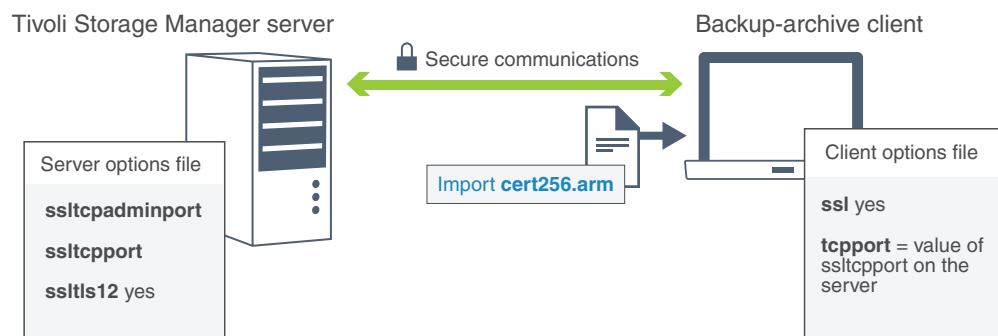
When you updated the server options file in "Setting options for the server" on page 51, the SSLTLS12 and SSLDISABLELEGACYTLS server options were set to restrict secure communications to use TLS 1.2. This setting prevents the use of previous TLS protocol levels, which are less secure.

Procedure

To configure the Tivoli Storage Manager server and clients for SSL/TLS, complete the following steps:

1. Create the key database file, `dsmcert.kdb`, on each client. Issue the following command in the `bin` directory on the client:

   ```
   gsk8capicmd_64 -keydb -create -populate
   -db dsmcert.kdb -pw password -stash
   ```

2. Change the default certificate in the `cert.kdb` key ring database file to the "TSM Server SelfSigned SHA Key" label. Issue the following command from the server instance directory:

   ```
   gsk8capicmd_64 -cert -setdefault -db cert.kdb
   -stashed -label "TSM Server SelfSigned SHA Key"
   ```

3. Manually transfer the Tivoli Storage Manager server `cert256.arm` file to the client computers.

   The `cert256.arm` is created in the server instance directory when the **SSLTCPPORT** server option is specified.

4. Specify the following options in the client options file:
   - Set the `ssl` option to **yes**.
   - Set the `tcpport` option value to match the value of the SSLTCPPORT that is set on the server.

Configuring the Operations Center

After you install the Operations Center, complete the following configuration steps to start managing your storage environment.

Before you begin

When you connect to the Operations Center for the first time, you must provide the following information:
- Connection information for the Tivoli Storage Manager server that you want to designate as a hub server
- Login credentials for an administrator ID that is defined for that Tivoli Storage Manager server

Procedure

1. Designate the hub server. In a web browser, enter the following address:

   ```
   https://hostname:secure_port/oc
   ```

   where:
   - *hostname* represents the name of the computer where the Operations Center is installed
   - *secure_port* represents the port number that the Operations Center uses for HTTPS communication on that computer

   For example, if your host name is *tsm.storage.mylocation.com* and you are using the default secure port for the Operations Center, which is 11090, the address is:

   ```
   https://tsm.storage.mylocation.com:11090/oc
   ```

   When you log in to the Operations Center for the first time, a wizard guides you through an initial configuration to set up a new administrator with system authority on the server.

2. Set up secure communications between the Operations Center and the hub server by configuring the Secure Sockets Layer (SSL) protocol.

   Follow the instructions in "Securing communications between the Operations Center and the hub server."

3. Optional: To receive a daily email report that summarizes system status, configure your email settings in the Operations Center.

   Follow the instructions in "Tracking system status by using daily email reports" on page 89.

Securing communications between the Operations Center and the hub server

To secure communications between the Operations Center and the hub server by using the Secure Sockets Layer (SSL) protocol, you must add the SSL certificate of the hub server to the truststore file of the Operations Center.

Before you begin

The truststore file of the Operations Center is a container for SSL certificates that the Operations Center can access. It contains the SSL certificate that the Operations Center uses for HTTPS communication with web browsers.

During the installation of the Operations Center, you create a password for the truststore file. To set up SSL communication between the Operations Center and the hub server, you must use the same password to add the SSL certificate of the hub server to the truststore file. If you do not remember this password, you can reset it.

Figure 2 illustrates the components for setting up SSL between the Operations Center and the hub server.
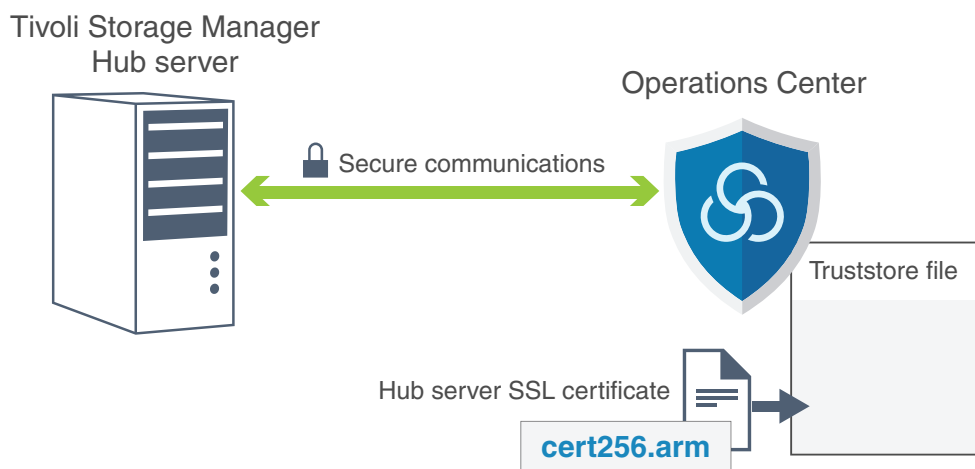

Figure 2. Configuring SSL between the Operations Center and the hub server

About this task

This procedure provides steps to implement secure communications by using self-signed certificates. To use certificate authority (CA) certificates, follow the instructions in Configuring SSL and TLS by using CA-signed certificates

(http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.admin/
t_ssl_config_ca.html).

Procedure

To set up SSL communication by using self-signed certificates, complete the
following steps.

1. Specify the `cert256.arm` certificate as the default certificate in the key database
   file of the hub server:

   a. Issue the following command from the hub server instance directory:
      ```
      gsk8capicmd_64 -cert -setdefault -db cert.kdb -stashed
            -label "TSM Server SelfSigned SHA Key"
      ```

   b. Restart the hub server so that it can receive the changes to the key
      database file.

   c. Verify that the `cert256.arm` certificate is set as the default. Issue the
      following command:
      ```
      gsk8capicmd_64 -cert -list -db cert.kdb -stashed
      ```

2. Stop the Operations Center web server.

3. Open the operating system command-line on the system where the Operations
   Center is installed, and change to the following directory:

   - **AIX** **Linux** `installation_dir/ui/jre/bin`

   - **Windows** `installation_dir\ui\jre\bin`

   Where *installation_dir* represents the directory in which the Operations Center
   is installed.

4. Open the IBM Key Management window by issuing the following command:
   ```
   ikeyman
   ```

5. Click **Key Database File > Open**.

6. Click **Browse**, and go to the following directory, where *installation_dir*
   represents the directory in which the Operations Center is installed:

   - **AIX** **Linux** `installation_dir/ui/Liberty/usr/servers/guiServer`

   - **Windows** `installation_dir\ui\Liberty\usr\servers\guiServer`

7. In the `guiServer` directory, select the `gui-truststore.jks` file.

8. Click **Open**, and click **OK**.

9. Enter the password for the truststore file, and click **OK**.

10. In the Key database content area of the IBM Key Management window, click
    the arrow, and select **Signer Certificates** from the list. Click **Add**.

11. In the Open window, click **Browse**, and go to the hub server instance
    directory:

    - **AIX** **Linux** `/opt/tivoli/tsm/server/bin`

    - **Windows** `c:\Program Files\Tivoli\TSM\server1`

    The directory contains the following SSL certificates:
    ```
    cert.arm
    cert256.arm
    ```

    If you cannot access the hub server instance directory from the Open window,
    complete the following steps:

a. Use FTP or another file-transfer method to copy the `cert256.arm` files from the hub server to the following directory on the computer where the Operations Center is installed:

- `AIX` `Linux` *installation_dir*/ui/Liberty/usr/servers/guiServer

- `Windows` *installation_dir*\ui\Liberty\usr\servers\guiServer

b. In the Open window, go to the `guiServer` directory.

12. Select the `cert256.arm` certificate as the SSL certificate.

13. Click **Open**, and click **OK**.

14. Enter a label for the certificate. For example, enter the name of the hub server.

15. Click **OK**. The SSL certificate of the hub server is added to the truststore file, and the label is displayed in the Key database content area of the IBM Key Management window.

16. Close the IBM Key Management window.

17. Start the Operations Center web server.

18. Complete the following steps in the login window of the configuration wizard:

a. In the **Connect to** field, enter the value of the **SSLTCPADMINPORT** server option as the port number.

b. Select **Use SSL**.

Registering your Tivoli Storage Manager license

To register a license for Tivoli Storage Manager, use the **REGISTER LICENSE** command.

About this task

Licenses are stored in enrollment certificate files, which contain licensing information for the product. The enrollment certificate files are on the installation media, and are placed on the server during installation. When you register the product, the licenses are stored in a `NODELOCK` file within the current directory.

Procedure

Register a Tivoli Storage Manager license by specifying the name of the enrollment certificate file that contains the license. To use the Operations Center command builder for this task, complete the following steps.

1. Open the Operations Center.
2. Open the Operations Center command builder by hovering over the globe icon and clicking **Command Builder**.
3. Issue the **REGISTER LICENSE** command. For example, to register a base Tivoli Storage Manager license, issue the following command:

   `register license file=tsmbasic.lic`

What to do next

Save the installation media that contains your enrollment certificate files. You might need to register your license again if, for example, one of the following conditions occur:

- The server is moved to a different computer.

- The NODELOCK file is corrupted. Tivoli Storage Manager stores license information in the NODELOCK file, which is in the directory from which the server is started.
- Linux If you change the processor chip that is associated with the server on which the Tivoli Storage Manager is installed.

Related reference:

➥ REGISTER LICENSE (Register a new license) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_license_register.html)

Configuring data deduplication

Create a directory-container storage pool and at least one directory to use inline data deduplication.

Before you begin

Use the storage pool directory information that you recorded in Chapter 4, "Planning worksheets," on page 11 for this task.

Procedure

1. Open the Operations Center.
2. On the Operations Center menu bar, hover over **Storage**.
3. From the list that is displayed, click **Storage Pools**.
4. Click the **+Storage Pool** button.
5. Complete the steps in the Add Storage Pool wizard:
   - To use inline data deduplication, select a **Directory** storage pool under Container-based storage.
   - When you configure directories for the directory-container storage pool, specify the directory paths that you created for storage during system setup.
6. After you configure the new directory-container storage pool, click **Close & View Policies** to update your management class and start using the storage pool.

Defining data retention rules for your business

After you create a directory-container storage pool for data deduplication, update the default server policy to use the new storage pool. The Add Storage Pool wizard opens the Services page in the Operations Center to complete this task.

Procedure

1. On the Services page of the Operations Center, select the STANDARD domain and click **Details**.
2. On the Summary page for the policy domain, click the **Policy Sets** tab. The Policy Sets page indicates the name of the active policy set and lists all of the management classes for that policy set.
3. Click the **Configure** toggle, and make the following changes:
   - Change the backup destination for the STANDARD management class to the directory-container storage pool.
   - Change the value for the Backups column to **No limit**.
   - Change the retention period. Set the Keep Extra Backups column to 30 days or more, depending on your business requirements.

4. Save your changes and click the **Configure** toggle again so that the policy set is no longer editable.

5. Activate the policy set by clicking **Activate**.

Related tasks:

"Specifying rules for backing up and archiving client data" on page 104

Defining schedules for server maintenance activities

Create schedules for each server maintenance task by using the **DEFINE SCHEDULE** command in the Operations Center command builder.

About this task

Schedule server maintenance operations to run after the client backup. You can control the timing of schedules for maintenance tasks by setting the start time in combination with the duration time for each operation.

The following example shows how you can schedule server maintenance processes in combination with the client backup schedule for a multisite disk solution.

Operation	Schedule
Client backup	Starts at 22:00.
Node replication	Starts at 08:00, or 10 hours after the beginning of the client backup.
Processing for database and disaster recovery files	• Database backup starts at 11:00, or 13 hours after the beginning of the client backup. This process runs until completion. • Device configuration information and volume history backup starts at 17:00, or 6 hours after the start of the database backup. • Volume history deletion starts at 20:00, or 9 hours after the start of the database backup.
Inventory expiration	Starts at 12:00, or 14 hours after the beginning of the client backup window. This process runs until completion.

Procedure

To run the server maintenance operations, set up schedules for each task by using the **DEFINE SCHEDULE** command.

You will create the schedule for replication separately, through the Operations Center, in a later step.

Task	Example command

1. Before you create the schedule for the database backup, you must define a device class for the backup operation.	Use the **DEFINE DEVCLASS** command to create a device class that is named DBBACK_FILEDEV:
	
	`define devclass dbback_filedev devtype=file`
	`  directory=`*`db_backup_directories`*
	
	where *db_backup_directories* is a list of the directories that you created for the database backup.
	
	AIX　**Linux**　For example, if your database backup directories are /tsminst1/TSMbkup00-03, issue the following command:
	
	`define devclass dbback_filedev devtype=file`
	`  directory=/tsminst1/TSMbkup00,`
	`  /tsminst1/TSMbkup01,/tsminst1/TSMbkup02,`
	`  /tsminst1/TSMbkup03"`
	
	Windows　For example, if your database backup directories are C:\tsminst1\TSMbkup00-03, issue the following command:
	
	`define devclass dbback_filedev devtype=file`
	`  directory="c:\tsminst1\TSMbkup00,`
	`  c:\tsminst1\TSMbkup01,c:\tsminst1\TSMbkup02,`
	`  c:\tsminst1\TSMbkup03"`

2. Set the device class for automatic database backups.	Use the **SET DBRECOVERY** to specify the device class that you created for the database backup: `set dbrecovery dbback_filedev`

3. Create schedules for the remaining operations by using the **DEFINE SCHEDULE** command for each task.	Create a schedule to run the **BACKUP DB** command. If you are configuring a small system, set the **COMPRESS** parameter to YES.
	
	For example, on a small system, you would issue the following command to create a backup schedule that uses the new device class:
	
	`define schedule DBBACKUP type=admin cmd="backup db`
	`  devclass=dbback_filedev type=full numstreams=3 wait=yes`
	`  compress=yes" active=yes desc="Back up the database."`
	`  startdate=today starttime=11:00:00 duration=45 durunits=minutes`

4. Back up the device configuration information.	Create a schedule to run the **BACKUP DEVCONFIG** command:
	
	`define schedule DEVCONFIGBKUP type=admin cmd="backup devconfig`
	`  filenames=devconfig.dat" active=yes desc="Backup the device`
	`  configuration file." startdate=today starttime=17:00:00`
	`  duration=45 durunits=minutes`

5. Back up the volume history.	Create a schedule to run the **BACKUP VOLHISTORY** command:
	
	`define schedule VOLHISTBKUP type=admin cmd="backup volhistory`
	`  filenames=volhist.dat" active=yes desc="Back up the volume`
	`  history." startdate=today starttime=17:00:00 duration=45`
	`  durunits=minutes`

6. Remove older versions of database backups, that are no longer required.	Create a schedule to run the **DELETE VOLHISTORY** command:
	
	`define schedule DELVOLHIST type=admin cmd="delete volhistory`
	`  type=dbb todate=today-6 totime=now" active=yes desc="Remove`
	`  old database backups." startdate=today starttime=20:00:00`
	`  duration=45 durunits=minutes`

Task	Example command

7. Remove objects that exceed their allowed retention.	Create a schedule to run the **EXPIRE INVENTORY** command.
	
	Set the **RESOURCE** parameter based on the system size that you are configuring:
	• Small systems: 10
	• Medium systems: 30
	• Large systems: 40
	
	For example, on a medium-sized system, issue the following command to create a schedule that is named EXPINVENTORY:
	
	`define schedule EXPINVENTORY type=admin cmd="expire inventory`
	`  wait=yes resource=30 duration=120" active=yes desc="Remove expired`
	`  objects." startdate=today starttime=12:00:00 duration=45`
	`  durunits=minutes`

What to do next

After you create schedules for the server maintenance tasks, you can view them in the Operations Center by completing the following steps:

1. On the Operations Center menu bar, hover over **Servers**.
2. Click **Maintenance**.

Related reference:

➡ DEFINE SCHEDULE (Define a schedule for an administrative command) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/ r_cmd_schedule_admin_define.html)

Defining client schedules

Use the Operations Center to create schedules for client operations.

Procedure

1. On the Operations Center menu bar, hover over **Clients**.
2. Click **Schedules**.
3. Click **+Schedule**.
4. Complete the steps in the Create Schedule wizard.
   • Set client backup schedules to start at 22:00, based on the server maintenance activities that you scheduled in "Defining schedules for server maintenance activities" on page 58.

Chapter 10. Installing and configuring clients

Following the successful setup of your server system, install, and configure client software to begin backing up data.

Procedure

To install the Tivoli Storage Manager backup-archive client, follow the installation instructions for your operating system.

* Installing the UNIX and Linux backup-archive clients (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/client/c_inst_baunix.html)
* Installing the Windows backup-archive client (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/client/t_inst_winclient.html)

What to do next

Register and assign your clients to schedules.

Registering and assigning clients to schedules

Add and register your clients through the Operations Center by using the Add Client wizard.

Procedure

1. On the Operations Center menu bar, click **Clients**.
2. In the Clients table, click **+Client**.
3. Complete the steps in the Add Client wizard:
   a. Optional: Specify that redundant data can be eliminated on both the client and server. In the Client-side deduplication area, select the **Enable** check box.
   b. In the Configuration window, copy the `TCPSERVERADDRESS`, `TCPPORT`, `NODENAME`, and `DEDUPLICATION` option values.

      Tip: Record the option values and keep them in a safe place. After you complete the client registration and install the software on the client node, use the values to configure the client.
   c. Follow the instructions in the wizard to specify the policy domain, schedule, and option set.
   d. Optional: Set how risks are displayed for the client by specifying the at-risk setting.
   e. Click **Add Client**.

Installing the client management service

Install the client management service for backup-archive clients that run on Linux and Windows operating systems. The client management service collects diagnostic information about IBM Tivoli Storage Manager backup-archive clients and makes the information available to the Operations Center for basic monitoring capability.

Procedure

Install the client management service on the same computer as the backup-archive client by completing the following steps:

1. Obtain the installation package for the client management service from the product DVD. Alternatively, you can download the installation package for the client management service from an IBM download site such as IBM Passport Advantage® or IBM Fix Central. Look for a file name that is similar to `version-TIV-TSMCMS-operating_system.bin`.
2. Create a directory on the client system that you want to manage, and copy the installation package there.
3. Extract the contents of the installation package file.
4. Run the installation batch file from the directory where you extracted the installation and associated files. This is the directory that you created in step 2.
5. To install the client management service, follow the instructions in the IBM Installation Manager wizard. If IBM Installation Manager is not already installed on the client system, you must select both IBM Installation Manager and IBM Tivoli Storage Manager Client Management Services.

Related tasks:

➡ Configuring the client management service for custom client installations (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.install/t_oc_inst_cms_customcfg.html)

Verifying the client management service is installed correctly

Before you use the client management service to collect diagnostic information about a backup-archive client, you can verify that the client management service is correctly installed and configured.

Procedure

1. Open a command line on the client system and issue the following command to view the configuration of the client management service: In a web browser, enter the following address, where:

   Linux

   `<client_install_dir>/cms/bin/CmsConfig.sh list`

   Windows

   `<client_install_dir>\cms\bin\CmsConfig.bat list`

   where *client_install_dir* is the directory where the backup-archive client is installed.

   The output is similar to the following: **Linux**

   ```
   List:
   server1.example.com:1500 HOSTNAME
       /opt/tivoli/tsm/client/ba/bin/dsmerror.log
           en_US MM/dd/yyyy HH:mm:ss UTF-8
   ```

```
List:
server1.example.com:1500 HOSTNAME
    C:/Program Files/Tivoli/TSM/baclient/dsmerror.log
        en_US MM/dd/yyyy HH:mm:ss Windows-1252
```

2. If the client management service is correctly installed and configured, the output displays the location of the error log file. The output text is extracted from the following configuration file:

   - Linux `client_install_dir/cms/Liberty/usr/servers/cmsServer/client-configuration.xml`

   - Windows `client_install_dir\cms\Liberty\usr\servers\cmsServer\client-configuration.xml`

   If the output does not contain any entries, you must configure the `client-configuration.xml` file. For instructions about how to configure this file, see Configuring the client management service for custom client installations (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.install/t_oc_inst_cms_customcfg.html).

Configuring the Operations Center to use the client management service

If you did not use the default configuration for the client management service, you must configure the Operations Center to access the client management service.

Before you begin

Ensure that the client management service is installed and started on the client system. Verify whether the default configuration is used. The default configuration is not used if either of the following conditions is met:

- The client management service does not use the default port number, 9028.
- The backup-archive client is not accessed by the same IP address as the client system where the backup-archive client is installed. For example, a different IP address might be used in the following situations:
  - The computer system has two network cards. The backup-archive client is configured to communicate on one network, while the client management service communicates on the other network.
  - The client system is configured with the Dynamic Host Configuration Protocol (DHCP). As a result, the client system is dynamically assigned an IP address, which is saved on the Tivoli Storage Manager server during the previous backup-archive client operation. When the client system is restarted, the client system might be assigned a different IP address. To ensure that the Operations Center can always find the client system, you specify a fully qualified domain name.

Procedure

To configure the Operations Center to use the client management service, complete the following steps:

1. On the Clients page of the Operations Center, select the client.
2. Click **Details** > **Properties**.
3. In the Remote diagnostics URL field in the General section, specify the URL for the client management service on the client system. The address must start with `https`. The following table shows examples of the remote diagnostics URL.

Type of URL	Example
With DNS host name and default port, 9028	`https://server.example.com`
With DNS host name and non-default port	`https://server.example.com:1599`
With IP address and non-default port	`https://192.0.2.0:1599`

What to do next

You can access client diagnostic information such as client log files from the **Diagnosis** tab in the Operations Center.

Chapter 11. Configuring the second server

After you complete the configuration for the first server in your system, configure the second server.

Procedure

To configure a second server that is the same as the first server, complete the instructions in the following sections:

- Chapter 7, "Setting up the system," on page 29
- Chapter 8, "Installing the Tivoli Storage Manager server and Operations Center," on page 45

  Only one server in the multisite disk solution is configured as the hub server, so you do not need to install the Operations Center on the second server. When you select the installation packages to install on the second server, do not select the Operations Center.
- Chapter 9, "Configuring the Tivoli Storage Manager server and the Operations Center," on page 49

  Skip the tasks for configuring the Operations Center.
- Chapter 10, "Installing and configuring clients," on page 61

What to do next

Complete the configuration by configuring SSL for the hub and spoke servers, adding the second server as a spoke to the hub server, and enabling replication.

Configuring SSL communications between the hub server and a spoke server

To secure communications between the hub server and a spoke server by using the Secure Sockets Layer (SSL) protocol, you must define the SSL certificate of the spoke server to the hub server. You must also configure the Tivoli Storage Manager to monitor the spoke server.

Procedure

1. To ensure that SSL ports are correctly set on the hub server and each spoke server, complete the following steps:

   a. From a Tivoli Storage Manager command line, issue the following command to each server:

   ```
   QUERY OPTION SSL*
   ```

   The results include the server options that are shown in the following example:

   ```
   Server Option        Option Setting
   ------------------------------------
   SSLTCPPort           3700
   SSLTCPADMINPort      3800
   SSLTLS12             Yes
   SSLFIPSMODE          No
   ```

   b. Ensure the following:

- The **SSLTCPPORT** and **SSLTCPADMINPORT** options have values in the Option Setting column.
- The **SSLTLS12** option is set to YES so that the Transport Layer Security (TLS) protocol version 1.2 is used for communication.

To update the values of these options, edit the `dsmserv.opt` file of the respective server, and restart that server.

2. On the spoke server, change to the directory of the spoke server instance.
3. Specify the required `cert256.arm` certificate as the default certificate in the key database file of the spoke server. Issue the following command:

```
gsk8capicmd_64 -cert -setdefault -db cert.kdb -stashed
 -label "TSM Server SelfSigned SHA Key"
```

4. Verify the certificates in the key database file of the spoke server. Issue the following command:

```
gsk8capicmd_64 -cert -list -db cert.kdb -stashed
```

The command generates output that is similar to the following example:

```
Certificates found
* default, - personal, ! trusted
!       Entrust.net Secure Server Certification Authority
!       Entrust.net Certification Authority (2048)
!       Entrust.net Client Certification Authority
!       Entrust.net Global Client Certification Authority
!       Entrust.net Global Secure Server Certification Authority
!       VeriSign Class 1 Public Primary Certification Authority
!       VeriSign Class 2 Public Primary Certification Authority
!       VeriSign Class 3 Public Primary Certification Authority
!       VeriSign Class 1 Public Primary Certification Authority - G2
!       VeriSign Class 2 Public Primary Certification Authority - G2
!       VeriSign Class 3 Public Primary Certification Authority - G2
!       VeriSign Class 4 Public Primary Certification Authority - G2
!       VeriSign Class 1 Public Primary Certification Authority - G3
!       VeriSign Class 2 Public Primary Certification Authority - G3
!       VeriSign Class 3 Public Primary Certification Authority - G3
!       VeriSign Class 3 Public Primary Certification Authority - G5
!       VeriSign Class 4 Public Primary Certification Authority - G3
!       VeriSign Class 3 Secure Server CA
!       Thawte Primary Root CA
!       Thawte Primary Root CA - G2 ECC
!       Thawte Server CA
!       Thawte Premium Server CA
!       Thawte Personal Basic CA
!       Thawte Personal Freemail CA
!       Thawte Personal Premium CA
-       TSM Server SelfSigned Key
*-      TSM Server SelfSigned SHA Key
```

5. Securely transfer the `cert256.arm` file of the spoke server to the hub server.
6. On the hub server, change to the directory of the hub server instance.
7. Define the spoke server SSL certificate to the hub server. Issue the following command from the hub server instance directory, where *spoke_servername* is the name of the spoke server, and *spoke_cert256.arm* is the file name of the spoke server SSL certificate:

```
gsk8capicmd_64 -cert -add -db cert.kdb -stashed -format ascii
 -label spoke_servername -file spoke_cert256.arm
```

The spoke server does not require the hub server SSL certificate for hub-to-spoke communication. However, other Tivoli Storage Manager server configurations that require cross-defined servers do require the spoke server to have the hub server SSL certificate.

> **Tip:** From each server, you can view the certificates in the key database file by issuing the following command:
>
> ```
> gsk8capicmd_64 -cert -list -db cert.kdb -stashed
> ```

8. Restart the hub server and the spoke server.

9. For the hub server, issue the **DEFINE SERVER** command, according to the following example:

   ```
   DEFINE SERVER spoke_servername HLA=spoke_address
    LLA=spoke_SSLTCPADMINPort SERVERPA=spoke_serverpassword SSL=YES
   ```

10. On the Operations Center menu bar, click **Servers**.

    In the table on the Servers page, the spoke server that you defined in 9 should have a status of "Unmonitored." Depending on the setting for the status refresh interval, you might not see the spoke server immediately.

11. Click the spoke server to highlight it, and in the table menu bar, click **Monitor Spoke**.

Related reference:

➡ DEFINE SERVER (Define a server for server-to-server communications) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_server_define.html)

➡ QUERY OPTION (Query server options) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_option_query.html)

Adding the second server as a spoke

After you configure both servers in your environment, add the second server as a spoke to the hub server.

Procedure

1. Open the Operations Center.
2. In the Operations Center menu bar, click **Servers**.
3. Complete one of the following steps:
   - Click the server to highlight it, and in the table menu bar, click **Monitor Spoke**.
   - If the server that you want to add is not shown in the table, click **+Spoke**.
4. Complete the steps in the spoke configuration wizard.

Enabling replication

To protect your data, enable node replication in addition to protecting your storage pools.

Procedure

To enable node replication for all of the clients that are registered to the source server, complete the following steps

1. Open the Operations Center.
2. On the Operations Center menu bar, hover over **Storage** and click **Replication**.
3. On the **Replication** page, click **+Server Pair**.
4. Complete the steps in the Add Server Pair wizard:
   - Set the source server as the first server that you configured for the multisite disk solution. The target server is the second server.

- Set the node replication schedule to start 10 hours after the client backup window, based on the server maintenance activities that you scheduled in "Defining schedules for server maintenance activities" on page 58.
- The wizard sets up storage pool protection schedules for you, based on the amount of data that you are protecting and when client replication is scheduled.

What to do next

If you plan to set up mutual replication between the two sites, run the Add Server Pair wizard again and set the second server as the source and the first server as the target.

Chapter 12. Completing the implementation

After the solution is configured and running, test backup operations and set up monitoring to ensure that everything runs smoothly.

Procedure

1. Test backup operations to verify that your data is protected in the way that you expect.

   a. On the Clients page of the Operations Center, select the clients that you want to back up, and click **Back Up**.

   b. On the Servers page of the Operations Center, select the server that contains the database that you want to back up. Click **Back Up**, and follow the instructions in the Back Up Server Database window.

   c. Verify that the backups completed successfully and there are no warning or error messages.

2. Set up monitoring for your solution by following the instructions in Part 3, "Monitoring a multisite disk solution," on page 71.

Part 3. Monitoring a multisite disk solution

After you implement a Tivoli Storage Manager multisite disk solution, monitor the solution to ensure that it operates correctly. By monitoring the solution on a daily and periodic basis, you can identify existing and potential issues. The information that you gather can be used to troubleshoot problems and optimize system performance.

About this task

The preferred way to monitor a solution is by using the Operations Center, which provides overall and detailed system status in a graphical user interface. In addition, you can configure the Operations Center to generate a daily email report that summarizes system status.

In some cases, you might want to use advanced monitoring tools to complete specific monitoring or troubleshooting tasks. To select and configure advanced monitoring tools, see Chapter 16, "Selecting, configuring, and using monitoring tools," on page 87.

Tip: If you plan to diagnose issues with backup-archive clients on Linux or Windows operating systems, install Tivoli Storage Manager client management services on each computer where a backup-archive client is installed. In this way, you can ensure that the **Diagnose** button is available in the Operations Center for diagnosing issues with backup-archive clients. To install the client management service, follow the instructions in "Installing the client management service" on page 62.

Procedure

1. Complete daily monitoring tasks. For instructions, see Chapter 13, "Daily monitoring checklist," on page 73.
2. Complete periodic monitoring tasks. For instructions, see Chapter 14, "Periodic monitoring checklist," on page 79.
3. To verify that your Tivoli Storage Manager solution complies with licensing requirements, follow the instructions in Chapter 15, "Verifying license compliance," on page 85.

What to do next

Resolve any issues that you detect. To resolve an issue by changing the configuration of your solution, follow the instructions in Part 4, "Managing operations," on page 93. The following additional resources are available:

- To resolve performance issues, see Performance (V7.1.1) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.perf.doc/c_performance.html).
- To resolve other types of issues, see Troubleshooting (V7.1.1) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.tsm.doc/t_tshoot_tsm.html).

Chapter 13. Daily monitoring checklist

To ensure that you are completing the daily monitoring tasks for your solution, review the daily monitoring checklist.

Complete the daily monitoring tasks from the Operations Center Overview page. You can access the Overview page by opening the Operations Center and clicking **Overviews**.

The following figure shows the location for completing each task.

Tip: To run administrative commands for advanced monitoring tasks, use the Operations Center command builder. The command builder provides a type-ahead function to guide you as you enter commands. To open the command builder, go to the Operations Center Overview page. On the menu bar, hover over the globe icon 🌐 and click **Command Builder**.

The following table lists the daily monitoring tasks and provides instructions for completing each task.

Table 15. Daily monitoring tasks

Task	Basic procedures	Advanced procedures and troubleshooting information
1 Determine whether clients are at risk of being unprotected due to failed or missed backup operations.	To verify whether clients are at risk, in the Clients area, look for an **At risk** notification. To view details, click the Clients area. If you installed the client management service on a backup-archive client, you can view and analyze the client error and schedule logs by completing the following steps: 1. In the Clients table, select the client and click **Details**. 2. To diagnose an issue, click **Diagnosis**.	For clients that do not have the client management service installed, access the client system to review the client error logs.
2 Determine whether client-related or server-related errors require attention.	To determine the severity of any reported alerts, in the Alerts area, hover over the columns.	To view additional information about alerts, click the Alerts area to open the Alerts page. Activity log messages are available at the bottom of the Alerts page.
3 Determine whether servers that are managed by the Operations Center are available to provide data protection services to clients.	1. To verify whether servers are at risk, in the Servers area, look for an **Unavailable** notification. 2. To view additional information, click the Servers area. 3. Select a server in the Servers table and click **Details**.	**Tip:** If you detect an issue that is related to server properties, update the server properties: 1. In the Servers table, select a server and click **Details**. 2. To update server properties, click **Properties**.

Table 15. Daily monitoring tasks (continued)

Task	Basic procedures	Advanced procedures and troubleshooting information
4 Determine whether sufficient space is available for the server inventory, which consists of the server database, active log, and archive log.	1. Click the Servers area. 2. In the Status column of the table, view the status of the server and resolve any issues: • **Normal** ✓ Sufficient space is available for the server database, active log, and archive log. • **Critical** ✗ Insufficient space is available for the server database, active log, or archive log. You must add space immediately, or the data protection services that are provided by the server will be interrupted. • **Warning** ⚠ The server database, active log, or archive log is running out of space. If this condition persists, you must add space. • **Unavailable** ⟲ Status cannot be obtained. Ensure that the server is running, and that there are no network issues. This status is also shown if the monitoring administrator ID is locked or otherwise unavailable on the server. This ID is named IBM-OC-hub_server_name. • **Unmonitored** ◈ Unmonitored servers are defined to the hub server, but are not configured for management by the Operations Center. To configure an unmonitored server, select the server, and click **Monitor Spoke**.	You can also look for related alerts on the Alerts page. For additional instructions about troubleshooting, see Resolving server problems (V7.1.1) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.tshoot.doc/r_pdg_rsolvsrvrprobs.html).

Table 15. Daily monitoring tasks (continued)

Task	Basic procedures	Advanced procedures and troubleshooting information
5 Verify server database backup operations.	To determine when a server was most recently backed up, complete the following steps: 1. Click the Servers area. 2. In the Servers table, review the Last Database Backup column.	To obtain more detailed information about backup operations, complete the following steps: 1. In the Servers table, select a row and click **Details**. 2. In the DB Backup area, hover over the check marks to review information about backup operations. If a database was not backed up recently (for example, in the last 24 hours), you can start a backup operation: 1. On the Operations Center Overview page, click the Servers area. 2. In the table, select a server and click **Back Up**. To determine whether the server database is configured for automatic backup operations, complete the following steps: 1. On the menu bar, hover over the globe icon 🌐 and click **Command Builder**. 2. Issue the `QUERY DB` command: `query db f=d` 3. In the output, review the `Full Device Class Name` field. If a device class is specified, the server is configured for automatic database backups.
6 Monitor other server maintenance tasks. Server maintenance tasks can include running administrative command schedules, maintenance scripts, and related commands.	To search for information about processes that failed because of server issues, complete the following steps: 1. Click **Servers** > **Maintenance**. 2. To obtain the two-week history of a process, view the History column. 3. To obtain more information about a scheduled process, hover over the check box that is associated with the process.	For more information about monitoring processes and resolving issues, see the Operations Center online help.
7 Verify that the amount of data that was recently sent to and from servers is within the expected range.	• To obtain an overview of activity in the last 24 hours, view the Activity area. • To compare activity in the last 24 hours with activity in the previous 24 hours, review the figures in the Current® and Previous areas.	• If more data was sent to the server than you expected, determine which clients are backing up more data and investigate the cause. It is possible that client-side data deduplication is not working correctly. • If less data was sent to the server than you expected, investigate whether client backup operations are proceeding on schedule.

Table 15. Daily monitoring tasks (continued)

Task	Basic procedures	Advanced procedures and troubleshooting information
8 Verify that storage pools are available to back up client data.	1. If problems are indicated in the Storage & Data Availability area, click **Pools** to view the details: • If the **Critical** ❌ status is displayed, insufficient space is available in the storage pool, or its access status is unavailable. • If the **Warning** ⚠️ status is displayed, the storage pool is running out of space, or its access status is read-only. 2. To view the used, free, and total space for your selected storage pool, hover over the entries in the Capacity Used column.	To view the storage-pool capacity that was used over the past 2 weeks, select a row in the Storage Pools table and click **Details**.
9 Verify that storage devices are available for backup operations.	In the Storage & Data Availability area, click **Devices** to open the Storage Devices table. If a **Critical** ❌ or **Warning** ⚠️ status is displayed for any device, investigate the issue.	**Tip:** Devices might have a critical or warning status for the following reasons: • For DISK device classes, volumes might be offline or have a read-only access status. • For FILE device classes that are not shared, directories might be offline. Also, insufficient free space might be available for allocating scratch volumes.

Table 15. Daily monitoring tasks (continued)

Task	Basic procedures	Advanced procedures and troubleshooting information

10 Monitor node replication processes.	1. To obtain the overall status of node replication processes, view the Replication area on the Operations Center Overview page.
	
	2. To view information about each replicated server pair, click the Replication area.
	
	3. To view the amount of data that was replicated over the last 2 weeks and the speed of replication, select a server pair and click **Details**.
	
	4. To view replication information for a client, on the Operations Center Overview page, click **Clients**. View the information in the Replication Workload column.
		For advanced monitoring, view information about running and ended node replication processes by using commands:
		
		1. On the Operations Center Overview page, hover over the globe icon 🌐 and click **Command Builder**.
		
		2. Issue the **QUERY REPLICATION** command. For instructions, see QUERY REPLICATION (Query node replication processes) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_replication_query.html). If the replication operation was completed successfully, the `Total Files To Replicate` value matches the `Total Files Replicated` value.
		
		3. To view summary messages, issue the `SELECT * FROM SUMMARY WHERE ACTIVITY='REPLICATION'` command.
		
		To display messages that are related to a node replication process on a source replication server, complete the following steps:
		
		1. On the Operations Center Overview page, click **Servers**.
		
		2. Select the replication server and click **Details**:
		
		  • To view active tasks, click **Active Tasks**, select the task, and verify that the **Running** state is displayed. For details, view the related activity logs.
		
		  • To view completed tasks, click **Completed Tasks**, select the task, and ensure that the **Completed** status is displayed. For details, view the related activity logs.

Chapter 14. Periodic monitoring checklist

To help ensure that your solution operates correctly, complete the tasks in the periodic monitoring checklist. Schedule periodic tasks frequently enough so that you can detect potential issues before they become problematic.

Tip: To run administrative commands for advanced monitoring tasks, use the Operations Center command builder. The command builder provides a type-ahead function to guide you as you enter commands. To open the command builder, go to the Operations Center Overview page. On the menu bar, hover over the globe icon 🌐 and click **Command Builder**.

Table 16. Periodic monitoring tasks

Task	Basic procedures	Advanced procedures and troubleshooting
1. Monitor system performance.	Determine the length of time that is required for client backup operations: 1. On the Operations Center Overview page, click **Clients**. Find the server that is associated with the client. 2. Click **Servers**. Select the server and click **Details**. 3. To view the duration of completed tasks in the last 24 hours, click **Completed Tasks**. 4. To view the duration of tasks that were completed more than 24 hours ago, use the **QUERY ACTLOG** command. Follow the instructions in QUERY ACTLOG (Query the activity log) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_actlog_query.html). 5. If the duration of client backup operations is increasing and the reasons are not clear, investigate the cause. If you installed the client management service on a backup-archive client, you can diagnose performance issues for the backup-archive client by completing the following steps: 1. On the Operations Center Overview page, click **Clients**. 2. Select a backup-archive client and click **Details**. 3. To retrieve client logs, click **Diagnosis**.	For instructions about reducing the time that it takes for the client to back up data to the server, see Resolving common client performance problems (V7.1.1) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.perf.doc/c_bac_perf_opts.html). Look for performance bottlenecks. For instructions, see Identifying performance bottlenecks (V7.1.1) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.perf.doc/t_perf_pd.html). For information about identifying and resolving other performance issues, see Performance (V7.1.1) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.perf.doc/c_performance.html).

Table 16. Periodic monitoring tasks (continued)

Task	Basic procedures	Advanced procedures and troubleshooting
2. Determine the disk savings that are provided by data deduplication.	1. On the Operations Center Overview page, click **Pools**. 2. Select a pool and click **Quick Look**. 3. In the Data Deduplication area, view the Space saved row.	For advanced monitoring, to obtain detailed statistics about the data-deduplication process for a specific directory-container storage pool or cloud-container storage pool, complete the following steps: 1. On the Operations Center Overview page, hover over the globe icon 🌐 and click **Command Builder**. 2. Obtain a statistical report by issuing the **GENERATE DEDUPSTATS** command. Follow the instructions in GENERATE DEDUPSTATS (Generate data deduplication statistics for a directory-container storage pool) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_dedupstats_generate.html). 3. View the statistical report by issuing the **QUERY DEDUPSTATS** command. Follow the instructions in QUERY DEDUPSTATS (Query data deduplication statistics) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_dedupstats_query.html).
3. Verify that current backup files for device configuration and volume history information are saved.	Access your storage locations to ensure that the files are available. The preferred method is to save the backup files to two locations. To locate the volume history and device configuration files, complete the following steps: 1. On the Operations Center Overview page, hover over the globe icon 🌐 and click **Command Builder**. 2. To locate the volume history and device configuration files, issue the following commands: query option volhistory query option devconfig 3. In the output, review the Option Setting column to find the file locations. If a disaster occurs, both the volume history file and the device configuration file are required to restore the server database.	

Table 16. Periodic monitoring tasks (continued)

Task	Basic procedures	Advanced procedures and troubleshooting

4. Determine whether sufficient space is available for the instance directory file system.	Verify that at least 20% of free space is available in the instance directory file system. Take the action that is appropriate for your operating system:
	
	• **AIX** To view available space in the file system, on the operating system command line, issue the following command:
	
	`df -g instance_directory`
	
	where *instance_directory* specifies the instance directory.
	
	• **Linux** To view available space in the file system, on the operating system command line, issue the following command:
	
	`df -h instance_directory`
	
	where *instance_directory* specifies the instance directory.
	
	• **Windows** In the Windows Explorer program, right-click the file system and click **Properties**. View the capacity information.
	
	The preferred location of the instance directory depends on the operating system where the server is installed:
	
	• **AIX** **Linux**
	`/home/tsminst1/tsminst1`
	
	• **Windows** `C:\tsminst1`
	
	Tip: If you completed a planning worksheet, the location of the instance directory is recorded in the worksheet.	

5. Identify unexpected client activity.	To monitor client activity to determine whether data volumes exceed expected amounts, complete the following steps: 1. On the Operations Center Overview page, click the Clients area. 2. To view activity over the past two weeks, double-click any client. 3. To view the number of bytes sent to the client, click the **Properties** tab. 4. In the Last Session area, view the Sent to client row.	When you double-click a client in the Clients table, the Activity over 2 Weeks area displays the amount of data that the client sent to the server each day.

Table 16. Periodic monitoring tasks (continued)

Task	Basic procedures	Advanced procedures and troubleshooting
6. Monitor storage pool growth over time.	1. On the Operations Center Overview page, click the Pools area. 2. To view the capacity that was used over the last two weeks, select a pool and click **Details**.	**Tips:** • To control the duration that must elapse before all deduplicated extents are removed from a directory-container storage pool or cloud-container storage pool, after they are no longer referenced by the inventory, issue the **UPDATE STGPOOL** *pool_name* **REUSEDELAY=days** command. • To display information about updates to data extents in directory-container or cloud-container storage pools, issue the **QUERY EXTENTUPDATES** command. The command output can help you determine what data extents are no longer referenced and which ones are eligible to be deleted from the system.
7. Evaluate the timing of client schedules. Ensure that the start and end times of client schedules meet your business needs.	On the Operations Center Overview page, click **Clients** > **Schedules**. In the Schedules table, the Start column displays the configured start time for the scheduled operation. To see when the most recent operation was started, hover over the clock icon.	**Tip:** You can receive a warning message if a client operation runs longer than expected. Complete the following steps: 1. On the Operations Center Overview page, hover over **Clients** and click **Schedules**. 2. Select a schedule and click **Details**. 3. View the details of a schedule by clicking the blue arrow next to the row. 4. In the **Run time alert** field, specify the time when a warning message will be issued if the scheduled operation is not completed. 5. Click **Save**.
8. Evaluate the timing of maintenance tasks. Ensure that the start and end times of maintenance tasks meet your business needs.	On the Operations Center Overview page, click **Servers** > **Maintenance**. In the Maintenance table, review the information in the Last Run Time column. To see when the last maintenance task was started, hover over the clock icon.	**Tip:** If a maintenance task is running too long, change the start time or the maximum run time. Complete the following steps: 1. On the Operations Center Overview page, hover over the globe icon 🌐 and click **Command Builder**. 2. To change the start time or maximum run time for a task, issue the **UPDATE SCHEDULE** command. For instructions, see UPDATE SCHEDULE (Update a client schedule) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_schedule_client_update.html).

Related reference:

↱ QUERY ACTLOG (Query the activity log) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_actlog_query.html)

↱ UPDATE STGPOOL (Update a directory-container storage pool) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_stgpool_cont_update.html)

↱ QUERY EXTENTUPDATES (Query updated data extents) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/

r_cmd_extentupdates_query.html)

Chapter 15. Verifying license compliance

Verify that your Tivoli Storage Manager solution complies with the provisions of your licensing agreement. By verifying compliance regularly, you can track trends in data growth or processor value unit (PVU) usage. Use this information to plan for future license purchasing.

About this task

The method for verifying license compliance varies depending on the provisions of your Tivoli Storage Manager licensing agreement:

Front-end capacity licensing
> The front-end model determines license requirements based on the amount of primary data that is reported as being backed up by clients. Clients include applications, virtual machines, and systems.

Back-end capacity licensing
> The back-end model determines license requirements based on the terabytes of data that are stored in Tivoli Storage Manager primary storage pools and IBM Tivoli Storage Manager FastBack repositories.
>
> **Important:** To ensure the accuracy of front-end and back-end capacity estimates, install the latest version of the client software on each client node. The front-end and back-end capacity estimates that are provided by the Operations Center are for planning purposes only. You can compare the information against your current license terms, and can review the estimates regularly to track growth trends. For the approved methods of measuring front-end or back-end capacity, refer to your licensing agreement.

PVU licensing
> The PVU model is based on the use of PVUs by server devices.
>
> **Important:** The PVU calculations that are provided by Tivoli Storage Manager are considered estimates and are not legally binding. The PVU licensing information that is reported by Tivoli Storage Manager is not considered an acceptable substitute for the IBM License Metric Tool.

For the latest information about Tivoli Storage Manager licensing models, see the information about comparing bundles at the Tivoli Storage Manager family website (http://www.ibm.com/software/products/tivoli-storage-manager-family). If you have questions or concerns about licensing requirements, contact your Tivoli Storage Manager software provider.

Procedure

To monitor license compliance, complete the steps that correspond to the provisions of your licensing agreement.

Option	Description
Front-end model	1. On the Operations Center menu bar, hover over the globe icon ⊕ and click **Licensing**. The front-end capacity estimate is displayed on the Front-end Usage page. 2. If a value is displayed in the Not Reporting column, click the number to identify clients that did not report capacity usage. 3. To estimate capacity for clients that did not report capacity usage, go to the following FTP site, which provides measuring tools and instructions: ftp://public.dhe.ibm.com/storage/tivoli-storage-management/front_end_capacity_measurement_tools Follow the instructions in the licensing guide to measure front-end capacity by script. 4. Add the Operations Center estimate and any estimates that you obtained by using a script. 5. Verify that the estimated capacity complies with your licensing agreement.
Back-end model	1. On the Operations Center menu bar, click the globe icon ⊕ and click **Licensing**. 2. Click the **Back-end** tab. 3. Verify that the estimated amount of data complies with your licensing agreement.
PVU model	Follow the instructions in Assessing compliance with the PVU licensing model (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.admin/t_verify_compl_pvu_lic.html).

Chapter 16. Selecting, configuring, and using monitoring tools

Use the Operations Center to obtain an overview of system status and to drill down to more detailed information. In some cases, you might want to use advanced tools to collect specific monitoring information.

Procedure

Select and configure the monitoring tools that are appropriate for your solution.

Table 17. Monitoring tools

Tool type	Use cases	Links to more information
Operations Center	• Use the graphical user interface to review system status and diagnose issues. • Set up the Operations Center to send daily email summary reports. • Optional: Customize the alerts that are displayed in the Operations Center and set up email notifications about the alerts. For instructions, hover over the Operations Center menu bar, click the help icon, and click **Documentation**. • Optional: Monitor the storage environment remotely by viewing the Overview page in the web browser of a mobile device. For example, you can use the Apple Safari web browser on an Apple iPad device. Other mobile devices can also be used. **Tip:** If you install Tivoli Storage Manager client management services on a backup-archive client, you can use the Operations Center to obtain troubleshooting information for the backup-archive client. The client management service can be installed only on Linux or Windows operating systems.	• Part 3, "Monitoring a multisite disk solution," on page 71 • Chapter 8, "Installing the Tivoli Storage Manager server and Operations Center," on page 45 • "Configuring the Operations Center" on page 53 • "Installing the client management service" on page 62 • "Tracking system status by using daily email reports" on page 89

Table 17. Monitoring tools (continued)

Tool type	Use cases	Links to more information
Tivoli Storage Manager administrative commands	Review detailed information. Use the method that is appropriate for your solution: • To display messages that were generated by the server and client, use the **QUERY ACTLOG** command. **Tip:** You can run Tivoli Storage Manager administrative commands from the Operations Center command builder. • To monitor activities such as server migration and client logons, use the administrative client in console mode. Run the **dsmadmc -consolemode** command.	• Administrative commands (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_ref.html) • QUERY ACTLOG (Query the activity log) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_actlog_query.html) • Monitoring server activities from the administrative client (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/t_cmdline_monitor_srv.html) • Administrative client options (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmdline_adclient_options.html)
Event logging	Log server messages and most client messages as events to one or more repositories called receivers.	For instructions about using event logging to monitor a solution, see Logging Tivoli Storage Manager events to receivers (V7.1.1) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.srv.doc/t_events_logging_receivers.html).
SQL queries	Create and format customized queries of the server database. For example, you can query the SQL activity summary table to view statistics about client operations and server processes. To display all information in the summary table, issue the following command from the administrative client: `select * from summary`	Using SELECT commands (V7.1.1) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.srv.doc/t_cmd_select_issuing.html)
Operating system tools	Monitor and test system performance.	
Device-monitoring tools	Monitor devices for availability, capacity, and performance. For example, use IBM Tivoli Storage Productivity Center or tools that are included in device hardware packages.	To monitor overall device status by using Tivoli Storage Productivity Center, follow the instructions in Monitoring the status of resources (http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.7/com.ibm.tpc_V527.doc/fqz0_c_wg_monitoring_status.html). To monitor performance by using Tivoli Storage Productivity Center, follow the instructions in Monitoring the performance of resources (http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.7/com.ibm.tpc_V527.doc/fqz0_c_performance_monitoring.html).

Table 17. Monitoring tools (continued)

Tool type	Use cases	Links to more information
IBM Tivoli Monitoring for Tivoli Storage Manager	Monitor Tivoli Storage Manager servers and produce historical reports about server and client activities. **Tip:** The Operations Center is the preferred tool for monitoring. However, Tivoli Monitoring for Tivoli Storage Manager is useful for generating historical reports that are based on IBM Cognos® Business Intelligence technology. Tivoli Monitoring for Tivoli Storage Manager is included in the Tivoli Storage Manager software package.	Tivoli Monitoring for Tivoli Storage Manager (http://www.ibm.com/developerworks/community/wikis/home/wiki/Tivoli%20Storage%20Manager/page/Tivoli%20Monitoring%20for%20Tivoli%20Storage%20Manager)

Tracking system status by using daily email reports

Set up the Operations Center to generate daily email reports that summarize system status. Specify the recipients and the time of day when they will receive the report.

Before you begin

Before you set up daily email reports, ensure that the following requirements are met:

- A Simple Mail Transfer Protocol (SMTP) host server is available to send and receive reports by email. The server must be configured as an open mail relay. You must also ensure that the Tivoli Storage Manager server that sends email messages has access to the SMTP server. If the Operations Center is installed on a separate computer, that computer does not require access to the SMTP server.
- You must have system privilege for the Tivoli Storage Manager server to set up daily email reports.
- To specify the recipients, you can enter one or more email addresses or administrator IDs. If you plan to enter an administrator ID, the ID must be registered on the hub server and must have an email address associated with it. To specify an email address for an administrator, use the `EMAILADDRESS` parameter of the `UPDATE ADMIN` command.

About this task

The daily email report is based on a template and is delivered as an image that is similar to the Operations Center Overview page. The following figure shows the summary page of a daily email report.

Data Protection Report for OCSERVC-1_SERVER1

July 28, 2015 at 3:54 PM

Executive Summary

Backup success rate **95.4%** ↑ | Server availability **80%** ↓

At risk **13** of **282 Clients** | Unavailable **1** of **5 Servers**

Activity over 24 Hours

	Total	Applications	Virtual Machines	Systems
Current	2.2 GB	0.0 B	0.0 B	2.2 GB
Previous	2.2 GB	0.0 B	0.0 B	2.2 GB

Protected Clients **282** At Risk **13**

Applications	6	Virtual Machines	8	Systems	268
	At Risk 1 (17%)		At Risk 8 (100%)		At Risk 4 (1%)

Servers **5** Unavailable **1**

Inventory

DB2
- Databases ✓
- Active logs ✓
- Archive logs ✓

Storage & Data Availability

Pools		Storage Devices		Replication	
Disk ✓		Disk ✗		Workload 0.0 B	
Tape ✓		Tape ✓			
Copy ✗		Copy ✗			

Tip: In addition to the summary page, the daily email report includes an attachment. To find more detailed information, expand the sections in the attachment.

Procedure

To specify settings for a daily email report, complete the following steps:

1. On the Operations Center Overview page, hover over the globe icon 🌐 on the menu bar and click **Settings**.

2. Click the **Daily Reports** tab, complete the fields, and click **Save**.

Results

If you cannot view the image in a daily report, you might be using an email client that converts HTML to another format. For information about restrictions, see the Operations Center online help.

Related reference:

➯ UPDATE ADMIN (Update an administrator) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_admin_update.html)

Part 4. Managing operations

Use this information to manage operations for a multisite disk implementation, that includes a Tivoli Storage Manager server, that uses data deduplication for multiple locations.

Chapter 17. Managing the Operations Center

The Operations Center provides web and mobile access to status information about the Tivoli Storage Manager environment. You can use the Operations Center to monitor multiple servers and to complete some administrative tasks. The Operations Center also provides web access to the Tivoli Storage Manager command line.

Adding and removing spoke servers

In a multiple-server environment, you can connect the other servers, called *spoke servers*, to the hub server.

About this task

The spoke servers send alerts and status information to the hub server. The Operations Center shows you a consolidated view of alerts and status information for the hub server and any spoke servers.

Adding a spoke server

After you configure the hub server for the Operations Center, you can add one or more spoke servers to the hub server.

Procedure

1. In the Operations Center menu bar, click **Servers**. The Servers page opens.

   In the table on the Servers page, a server might have a status of "Unmonitored." This status means that although an administrator defined this server to the hub server by using the `DEFINE SERVER` command, the server is not yet configured as a spoke server.

2. Complete one of the following steps:
   - Click the server to highlight it, and in the table menu bar, click **Monitor Spoke**.
   - If the server that you want to add is not shown in the table, click **+ Spoke** in the table menu bar.

3. Provide the necessary information, and complete the steps in the spoke configuration wizard.

   Tip: If the event-record retention period of the Tivoli Storage Manager server is less than 14 days, the period is automatically reset to 14 days if you configure the server as a spoke server.

Removing a spoke server

You can remove a spoke server from the Operations Center.

About this task

You might need to remove a spoke server in the following situations, for example:
- You want to move the spoke server from one hub server to another hub server.
- You want to decommission the spoke server.

Procedure

To remove the spoke server from the group of servers that are managed by the hub server, complete the following steps:

1. From the Tivoli Storage Manager command line, issue the following command on the hub server:

   QUERY MONITORSETTINGS

2. From the output of the command, copy the name that is in the **Monitored Group** field.

3. Issue the following command on the hub server, where *group_name* represents the name of the monitored group, and *member_name* represents the name of the spoke server:

   DELETE GRPMEMBER *group_name member_name*

4. Optional: If you want to move the spoke server from one hub server to another hub server, do **not** complete this step. Otherwise, you can disable alerting and monitoring on the spoke server by issuing the following commands on the spoke server:

   SET STATUSMONITOR OFF
   SET ALERTMONITOR OFF

5. Optional: If the spoke server definition is used for other purposes, such as enterprise configuration, command routing, storing virtual volumes, or library management, do **not** complete this step. Otherwise, you can delete the spoke server definition on the hub server by issuing the following command on the hub server:

   DELETE SERVER *spoke_server_name*

Starting and stopping the web server

The web server of the Operations Center runs as a service and starts automatically. You might need to stop and start the web server, for example, to make configuration changes.

Procedure

Stop and start the web server.
- ░░░░ AIX ░░░░ From the */installation_dir*/ui/utils directory, where *installation_dir* represents the directory where the Operations Center is installed, issue the following commands:
  – To stop the server:
    ./stopserver.sh
  – To start the server:
    ./startserver.sh

- [Linux] Issue the following commands:
  - To stop the server:

    `service opscenter.rc stop`
  - To start the server:

    `service opscenter.rc start`
  - To restart the server:

    `service opscenter.rc restart`

  To determine whether the server is running, issue the following command:

  `service opscenter.rc status`

- [Windows] From the Services window, stop or start the **Tivoli Storage Manager Operations Center** service.

Restarting the initial configuration wizard

You might need to restart the initial configuration wizard of the Operations Center, for example, to make configuration changes.

Before you begin

To change the following settings, use the Settings page in the Operations Center rather than restarting the initial configuration wizard:
- The frequency at which status data is refreshed
- The duration that alerts remain active, inactive, or closed
- The conditions that indicate that clients are at risk

The Operations Center help includes more information about how to change these settings.

About this task

To restart the initial configuration wizard, you must delete a properties file that includes information about the hub server connection. However, any alerting, monitoring, at-risk, or multiserver settings that were configured for the hub server are not deleted. These settings are used as the default settings in the configuration wizard when the wizard restarts.

Procedure

1. Stop the Operations Center web server.
2. On the computer where the Operations Center is installed, go to the following directory, where *installation_dir* represents the directory in which the Operations Center is installed:

   - [AIX] [Linux] *installation_dir*/ui/Liberty/usr/servers/guiServer

   - [Windows] *installation_dir*\ui\Liberty\usr\servers\guiServer

   For example:

   - [AIX] [Linux] /opt/tivoli/tsm/ui/Liberty/usr/servers/guiServer

   - [Windows] c:\Program Files\Tivoli\TSM\ui\Liberty\usr\servers\guiServer

3. In the guiServer directory, delete the serverConnection.properties file.
4. Start the Operations Center web server.
5. Open the Operations Center.

6. Use the configuration wizard to reconfigure the Operations Center. Specify a new password for the monitoring administrator ID.

7. On any spoke servers that were previously connected to the hub server, update the password for the monitoring administrator ID by issuing the following command from the Tivoli Storage Manager command-line interface:

   UPDATE ADMIN IBM-OC-*hub_server_name new_password*

   Restriction: Do not change any other settings for this administrator ID. After you specify the initial password, this password is managed automatically by the Operations Center.

Changing the hub server

You can remove the hub server of the Operations Center, and configure another server as the hub server.

Procedure

1. Restart the initial configuration wizard of the Operations Center. As part of this procedure, you delete the existing hub server connection.
2. Use the wizard to configure the Operations Center to connect to the new hub server.

Related tasks:

Restoring the configuration to the preconfiguration state

If certain problems occur, you might want to restore the Operations Center configuration to the preconfigured state where the Tivoli Storage Manager servers are not defined as hub or spoke servers.

Procedure

To restore the configuration, complete the following steps:

1. Stop the Operations Center web server.
2. Unconfigure the hub server by completing the following steps:
   a. On the hub server, issue the following commands:

   ```
   SET MONITORINGADMIN ""
   SET MONITOREDSERVERGROUP ""
   SET STATUSMONITOR OFF
   SET ALERTMONITOR OFF
   REMOVE ADMIN IBM-OC-hub_server_name
   ```

   Tip: IBM-OC-*hub_server_name* represents the monitoring administrator ID that was automatically created when you initially configured the hub server.

   b. Reset the password for the hub server by issuing the following command on the hub server:

   ```
   SET SERVERPASSWORD ""
   ```

   Attention: Do not complete this step if the hub server is configured with other servers for other purposes, such as library sharing, exporting and importing of data, or node replication.

3. Unconfigure any spoke servers by completing the following steps:

a. On the hub server, to determine whether any spoke servers remain as members of the server group, issue the following command:

```
QUERY SERVERGROUP IBM-OC-hub_server_name
```

Tip: IBM-OC-*hub_server_name* represents the name of the monitored server group that was automatically created when you configured the first spoke server. This server group name is also the same as the monitoring administrator ID that was automatically created when you initially configured the hub server.

b. On the hub server, to delete spoke servers from the server group, issue the following command for each spoke server:

```
DELETE GRPMEMBER IBM-OC-hub_server_name spoke_server_name
```

c. After all spoke servers are deleted from the server group, issue the following commands on the hub server:

```
DELETE SERVERGROUP IBM-OC-hub_server_name
SET MONITOREDSERVERGROUP ""
```

d. On each spoke server, issue the following commands:

```
REMOVE ADMIN IBM-OC-hub_server_name
SETOPT PUSHSTATUS NO
SET ALERTMONITOR OFF
SET STATUSMONITOR OFF
```

e. On each spoke server, delete the definition of the hub server by issuing the following command:

```
DELETE SERVER hub_server_name
```

Attention: Do not complete this step if the definition is used for other purposes, such as library sharing, exporting and importing of data, or node replication.

f. On the hub server, delete the definition of each spoke server by issuing the following command:

```
DELETE SERVER spoke_server_name
```

Attention: Do not complete this step if the server definition is used for other purposes, such as library sharing, exporting and importing of data, or node replication.

4. Restore the default settings on each server by issuing the following commands:

```
SET STATUSREFRESHINTERVAL 5
SET ALERTUPDATEINTERVAL 10
SET ALERTACTIVEDURATION 480
SET ALERTINACTIVEDURATION 480
SET ALERTCLOSEDDURATION 60
SET STATUSATRISKINTERVAL TYPE=AP INTERVAL=24
SET STATUSATRISKINTERVAL TYPE=VM INTERVAL=24
SET STATUSATRISKINTERVAL TYPE=SY INTERVAL=24
SET STATUSSKIPASFAILURE YES TYPE=ALL
```

5. Restart the initial configuration wizard of the Operations Center.

Related tasks:

Chapter 18. Protecting applications, virtual machines, and systems

A Tivoli Storage Manager server protects data for clients, which can include applications, virtual machines, and systems. To start protecting client data, register the client node with the server and select a backup schedule to protect the client data.

Adding clients

After you implement a Tivoli Storage Manager solution, you can expand the solution by adding clients.

About this task

The procedure describes basic steps for adding a client. For more specific instructions about configuring clients, see the documentation for the product that you install on the client node. You can have the following types of client nodes:

Application client nodes
Application client nodes include email servers, databases, and other applications. For example, any of the following applications can be an application client node:
- FlashCopy Manager
- Tivoli Storage Manager for Databases
- IBM Tivoli Storage Manager for Enterprise Resource Planning
- IBM Tivoli Storage Manager for Mail
- IBM Tivoli Storage Manager for Virtual Environments

System client nodes
System client nodes include workstations, network-attached storage (NAS) file servers, and API clients.

Virtual machine client nodes
Virtual machine client nodes consist of an individual guest host within a hypervisor. Each virtual machine is represented as a Tivoli Storage Manager file space.

Procedure

To add a client, complete the following steps:
1. Select the software to install on the client node and plan the installation. Follow the instructions in "Selecting the software and planning the installation" on page 102.
2. Specify how to back up and archive client data. Follow the instructions in "Specifying rules for backing up and archiving client data" on page 104.
3. Specify when to back up and archive client data. Follow the instructions in "Scheduling backup and archive operations" on page 107.
4. To allow the client to connect to the server, register the client. Follow the instructions in "Registering clients" on page 108.

5. To start protecting a client node, install and configure the selected software on the client node. Follow the instructions in "Installing and configuring clients" on page 109.

Selecting the software and planning the installation

Different types of data require different types of protection. Identify the type of data that you must protect and select the appropriate software.

About this task

The preferred practice is to install the Tivoli Storage Manager backup-archive client on all client nodes so that you can configure and start the Tivoli Storage Manager client acceptor daemon (CAD) on the client node. The CAD is designed to efficiently run scheduled operations.

The CAD runs schedules for the following products: the backup-archive client, Tivoli Storage Manager for Databases, Tivoli Storage Manager for Enterprise Resource Planning, Tivoli Storage Manager for Mail, and Tivoli Storage Manager for Virtual Environments. If you install a product for which the CAD does not run schedules, you must follow the configuration instructions in the product documentation to ensure that scheduled operations can occur.

Procedure

Based on your goal, select the products to install and review the installation instructions.

Tip: If you install the client software now, you must also complete the client configuration tasks that are described in "Installing and configuring clients" on page 109 before you can use the client.

Goal	Product and description	Installation instructions
Protect a file server or workstation	The Tivoli Storage Manager backup-archive client backs up and archives files and directories from file servers and workstations to storage. You can also restore and retrieve backup versions and archived copies of files.	• Backup-archive client requirements (http://www.ibm.com/support/ knowledgecenter/SSGSG7_7.1.3/client/ c_inst_env.html) • Installing the UNIX and Linux backup-archive clients (http://www.ibm.com/support/ knowledgecenter/SSGSG7_7.1.3/client/ c_inst_baunix.html) • Installing the Windows backup-archive client (http://www.ibm.com/support/ knowledgecenter/SSGSG7_7.1.3/client/ t_inst_winclient.html)

Goal	Product and description	Installation instructions
Protect applications with snapshot backup and restore capabilities	FlashCopy Manager protects data with integrated, application-aware snapshot backup and restore capabilities. You can protect data that is stored by IBM DB2 database software and SAP, Oracle, Microsoft Exchange, and Microsoft SQL Server applications.	• Installing and upgrading IBM Tivoli Storage FlashCopy Manager for UNIX and Linux (http://www.ibm.com/support/knowledgecenter/SS36V9_4.1.3/fcm.unx/t_protect_fcmul.html) • Installing and upgrading Tivoli Storage FlashCopy Manager for VMware (http://www.ibm.com/support/knowledgecenter/SS36V9_4.1.3/fcm.vmware/t_fvm_ins_ovr.html) • Installing and upgrading IBM Tivoli Storage FlashCopy Manager for Windows (http://www.ibm.com/support/knowledgecenter/SS36V9_4.1.3/fcm.win/c_fcm_installing.html)
Protect an email application on an IBM Domino server	Tivoli Storage Manager for Mail: Data Protection for IBM Domino® automates data protection so that backups are completed without shutting down IBM Domino servers.	• Installation of Data Protection for IBM Domino on a UNIX, AIX, or Linux system (http://www.ibm.com/support/knowledgecenter/SSTG2D_7.1.0/com.ibm.itsm.mail.dom.unx.doc/concept_dominstall.html) • Installation of Data Protection for IBM Domino on a Windows system (http://www.ibm.com/support/knowledgecenter/SSTG2D_7.1.0/com.ibm.itsm.mail.dom.win.doc/concept_dominstall.html)
Protect an email application on a Microsoft Exchange server	Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server automates data protection so that backups are completed without shutting down Microsoft Exchange servers.	Installing, upgrading, and migrating IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server (http://www.ibm.com/support/knowledgecenter/SSTG2D_7.1.3/mail.exc/t_dpexc_install_ovr.html)
Protect an IBM DB2 database	The application programming interface (API) of the Tivoli Storage Manager backup-archive client can be used to back up DB2 data to a Tivoli Storage Manager server.	Installing the Tivoli Storage Manager backup-archive clients (UNIX, Linux, and Windows) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/client/c_inst.html)
Protect an IBM Informix database	The API of the Tivoli Storage Manager backup-archive client can be used to back up Informix data to a Tivoli Storage Manager server.	Installing the Tivoli Storage Manager backup-archive clients (UNIX, Linux, and Windows) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/client/c_inst.html)
Protect a Microsoft SQL database	Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server protects Microsoft SQL data.	Installing Data Protection for SQL Server on Windows Server Core (http://www.ibm.com/support/knowledgecenter/SSTFZR_7.1.3/db.sql/t_dpfcm_ins_srvcore_sql.html)
Protect an Oracle database	Tivoli Storage Manager for Databases: Data Protection for Oracle protects Oracle data.	Data Protection for Oracle installation (http://www.ibm.com/support/knowledgecenter/SSTFZR_7.1.3/db.orc/t_dporc_installing.html)

Goal	Product and description	Installation instructions
Protect an SAP environment	Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP provides protection that is customized for SAP environments. The product is designed to improve the availability of SAP database servers and reduce administration workload.	• Installing Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP for DB2 (http://www.ibm.com/support/knowledgecenter/SSZHVN_7.1.3/erp.db2/t_dperp_d_install_ovr.html) • Installing Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP for Oracle (http://www.ibm.com/support/knowledgecenter/SSZHVN_7.1.3/erp.orc/t_dperp_o_install_ovr.html)
Protect a virtual machine	Tivoli Storage Manager for Virtual Environments provides protection that is tailored for Microsoft Hyper-V and VMware virtual environments. You can use Tivoli Storage Manager for Virtual Environments to create incremental forever backups that are stored on a centralized Tivoli Storage Manager server, create backup policies, and restore virtual machines or individual files. Alternatively, use the Tivoli Storage Manager backup-archive client to back up and restore a full VMware or Microsoft Hyper-V virtual machine. You can also back up and restore files or directories from a VMware virtual machine.	• Installing Data Protection for Microsoft Hyper-V (http://www.ibm.com/support/knowledgecenter/SS8TDQ_7.1.3/ve.hv/c_ve_hv_inst_plan.html) • Installing and upgrading Data Protection for VMware (http://www.ibm.com/support/knowledgecenter/SS8TDQ_7.1.3/ve.inst/t_ve_installing.html) • Installing the Tivoli Storage Manager backup-archive clients (UNIX, Linux, and Windows) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/client/c_inst.html)

Tip: To use the client for space management, you can install Tivoli Storage Manager for Space Management or IBM Tivoli Storage Manager HSM for Windows. For product descriptions and installation instructions, see Tivoli Storage Manager for Space Management (http://www.ibm.com/support/knowledgecenter/SSSR2R_7.1.3) and Tivoli Storage Manager HSM for Windows (http://www.ibm.com/support/knowledgecenter/SSATMW_7.1.3).

Specifying rules for backing up and archiving client data

Before you add a client, ensure that appropriate rules are specified for backing up and archiving the client data. During the client registration process, you assign the client node to a policy domain, which has the rules that control how and when client data is stored.

Before you begin

Determine how to proceed:
• If you are familiar with the policies that are configured for your solution and you know that they do not require changes, continue with "Scheduling backup and archive operations" on page 107.
• If you are not familiar with the policies, follow the steps in this procedure.

About this task

Policies affect how much data is stored over time, and how long data is retained and available for clients to restore. To meet objectives for data protection, you can update the default policy and create your own policies. A policy includes the following rules:

- How and when files are backed up and archived to server storage
- The number of copies of a file and the length of time copies are kept in server storage

During the client registration process, you assign a client to a *policy domain*. The policy for a specific client is determined by the rules in the policy domain to which the client is assigned. In the policy domain, the rules that are in effect are in the active *policy set*.

When a client backs up or archives a file, the file is bound to a management class in the active policy set of the policy domain. A *management class* is the key set of rules for managing client data. The backup and archive operations on the client use the settings in the default management class of the policy domain unless you further customize policy. A policy can be customized by defining more management classes and assigning their use through client options.

Client options can be specified in a local, editable file on the client system and in a client option set on the server. The options in the client option set on the server can override or add to the options in the local client option file.

Procedure

1. Review the policies that are configured for your solution by following the instructions in "Viewing policies."
2. If you need to make minor changes to meet data retention requirements, follow the instructions in "Editing policies" on page 106.
3. Optional: If you need to create policy domains or make extensive changes to policies to meet data retention requirements, see Customizing policies (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.admin/t_mplmntpol_getstrted.html).

Viewing policies

View policies to determine whether they must be edited to meet your requirements.

Procedure

1. To view the active policy set for a policy domain, complete the following steps:
   a. On the Services page of the Operations Center, select a policy domain and click **Details**.
   b. On the Summary page for the policy domain, click the **Policy Sets** tab.
2. To view inactive policy sets for a policy domain, complete the following steps:
   a. On the Policy Sets page, click the **Configure** toggle. You can now view and edit the policy sets that are inactive.
   b. Scroll through the inactive policy sets by using the forward and back arrows. When you view an inactive policy set, the settings that differentiate the inactive policy set from the active policy set are highlighted.
   c. Click the **Configure** toggle. The policy sets are no longer editable.

Editing policies

To change the rules that apply to a policy domain, edit the active policy set for the policy domain. You can also activate a different policy set for a domain.

Before you begin

Changes to policy can affect data retention. Ensure that you continue to back up data that is essential to your organization so that you can restore that data if a disaster occurs. Also, ensure that your system has sufficient storage space for planned backup operations.

About this task

You edit a policy set by changing one or more management classes within the policy set. If you edit the active policy set, the changes are not available to clients unless you reactivate the policy set. To make the edited policy set available to clients, activate the policy set.

Although you can define multiple policy sets for a policy domain, only one policy set can be active. When you activate a different policy set, it replaces the currently active policy set.

To learn about preferred practices for defining policies, see Customizing policies (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.admin/t_mplmntpol_getstrted.html).

Procedure

1. On the Services page of the Operations Center, select a policy domain and click **Details**.
2. On the Summary page for the policy domain, click the **Policy Sets** tab.
   The Policy Sets page indicates the name of the active policy set and lists all of the management classes for that policy set.
3. Click the **Configure** toggle. The policy set is editable.
4. Optional: To edit a policy set that is not active, click the forward and back arrows to locate the policy set.
5. Edit the policy set by completing any of the following actions:

Option	Description
Add a management class	1. In the Policy Sets table, click **+Management Class**. 2. To specify the rules for backing up and archiving data, complete the fields in the Add Management Class window. 3. To make the management class the default management class, select the **Make default** check box. 4. Click **Add**.
Delete a management class	In the Management Class column, click **-**. **Tip:** To delete the default management class, you must first assign a different management class as the default.

Option	Description
Make a management class the default management class	In the Default column for the management class, click the radio button. **Tip:** The default management class manages client files when another management class is not assigned to, or appropriate for managing, a file. To ensure that clients can always back up and archive files, choose a default management class that contains rules for both backing up and archiving files.
Modify a management class	To change the properties of a management class, update the fields in the table.

6. Click **Save**.

   Attention: When you activate a new policy set, data might be lost. Data that is protected under one policy set might not be protected under another policy set. Therefore, before you activate a policy set, ensure that the differences between the previous policy set and the new policy set do not cause data to be lost.

7. Click **Activate**. A summary of the differences between the active policy set and the new policy set is displayed. Ensure that the changes in the new policy set are consistent with your data retention requirements by completing the following steps:

   a. Review the differences between corresponding management classes in the two policy sets, and consider the consequences for client files. Client files that are bound to management classes in the active policy set will be bound to the management classes with the same names in the new policy set.

   b. Identify management classes in the active policy set that do not have counterparts in the new policy set, and consider the consequences for client files. Client files that are bound to these management classes will be managed by the default management class in the new policy set.

   c. If the changes to be implemented by the policy set are acceptable, select the **I understand that these updates can cause data loss** check box and click **Activate**.

Scheduling backup and archive operations

Before you register a new client with the server, ensure that a schedule is available to specify when backup and archive operations take place. During the registration process, you assign a schedule to the client.

Before you begin

Determine how to proceed:

- If you are familiar with the schedules that are configured for the solution and you know that they do not require modification, continue with "Registering clients" on page 108.

- If you are not familiar with the schedules or the schedules require modification, follow the steps in this procedure.

About this task

Typically, backup operations for all clients must be completed daily. Carefully schedule client and server workloads to achieve the best performance for your storage environment. To avoid the overlap of client and server operations, consider scheduling client backup and archive operations so that they run at night. If client

and server operations overlap or are not given enough time and resources to be processed, you might experience decreased system performance, failed operations, and other issues.

Procedure

1. Review available schedules by hovering over **Clients** on the Operations Center menu bar. Click **Schedules**.
2. Optional: Modify or create a schedule by completing the following steps:

Option	Description
Modify a schedule	1. In the Schedules view, select the schedule and click **Details**. 2. On the Schedule Details page, view details by clicking the blue arrows at the beginning of the rows. 3. Modify the settings in the schedule, and click **Save**.
Create a schedule	In the Schedules view, click **+Schedule** and complete the steps to create a schedule.

3. Optional: To configure schedule settings that are not visible in the Operations Center, use a server command. For example, you might want to schedule a client operation that backs up a specific directory and assigns it to a management class other than the default.
   a. On the Overview page of the Operations Center, hover over the globe icon 🌐 and click **Command Builder**.
   b. Issue the `DEFINE SCHEDULE` command to create a schedule or the `UPDATE SCHEDULE` command to modify a schedule. For details about the commands, see DEFINE SCHEDULE (Define a client schedule) (http://www.ibm.com/ support/knowledgecenter/SSGSG7_7.1.3/srv.reference/ r_cmd_schedule_client_define.html) or UPDATE SCHEDULE (Update a client schedule) (http://www.ibm.com/support/knowledgecenter/ SSGSG7_7.1.3/srv.reference/r_cmd_schedule_client_update.html).

Related tasks:

➡ Tuning the schedule for daily operations (V7.1.1) (http://www.ibm.com/ support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.perf.doc/ t_srv_tuning_daily.html)

Registering clients

Register a client to ensure that the client can connect to the Tivoli Storage Manager server, and the server can protect client data. Use the Add Client wizard in the Operations Center to register the client.

Procedure

1. On the Operations Center menu bar, click **Clients**.
2. In the Clients table, click **+Client**.
3. Complete the steps in the Add Client wizard:
   a. Optional: Specify that redundant data can be eliminated on both the client and server. In the Client-side deduplication area, select the **Enable** check box.

b. In the Configuration window, copy the **TCPSERVERADDRESS**, **TCPPORT**, **NODENAME**, and **DEDUPLICATION** option values.

   Tip: Record the option values and keep them in a safe place. After you complete the client registration and install the software on the client node, use the values to configure the client.

c. Follow the instructions in the wizard to specify the policy domain, schedule, and option set.

d. Optional: Set how risks are displayed for the client by specifying the at-risk setting.

e. Click **Add Client**.

Installing and configuring clients

To start protecting a client node, you must install and configure the selected software.

Procedure

If you already installed the software, start at step 2 on page 110.

1. Take one of the following actions:
   - To install software on an application or client node, follow the instructions.

Software	Link to instructions
Tivoli Storage Manager backup-archive client	• Installing the UNIX and Linux backup-archive clients (http://www.ibm.com/support/knowledgecenter/ SSGSG7_7.1.3/client/c_inst_baunix.html) • Installing the Windows backup-archive client (http://www.ibm.com/support/knowledgecenter/ SSGSG7_7.1.3/client/t_inst_winclient.html)
Tivoli Storage Manager for Databases	• Data Protection for Oracle installation (http://www.ibm.com/support/knowledgecenter/ SSTFZR_7.1.3/db.orc/t_dporc_installing.html) • Installing Data Protection for SQL Server on Windows Server Core (http://www.ibm.com/support/ knowledgecenter/SSTFZR_7.1.3/db.sql/ t_dpfcm_ins_srvcore_sql.html)
Tivoli Storage Manager for Mail	• Installation of Data Protection for IBM Domino on a UNIX, AIX, or Linux system (http://www.ibm.com/ support/knowledgecenter/SSTG2D_7.1.0/ com.ibm.itsm.mail.dom.unx.doc/concept_dominstall.html) • Installation of Data Protection for IBM Domino on a Windows system (http://www.ibm.com/support/ knowledgecenter/SSTG2D_7.1.0/ com.ibm.itsm.mail.dom.win.doc/concept_dominstall.html) • Installing, upgrading, and migrating IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server (http://www.ibm.com/support/knowledgecenter/ SSTG2D_7.1.3/mail.exc/t_dpexc_install_ovr.html)

Software	Link to instructions
Tivoli Storage FlashCopy Manager	• Installing and upgrading IBM Tivoli Storage FlashCopy Manager for UNIX and Linux (http://www.ibm.com/support/knowledgecenter/SS36V9_4.1.3/fcm.unx/t_protect_fcmul.html) • Installing and upgrading Tivoli Storage FlashCopy Manager for VMware (http://www.ibm.com/support/knowledgecenter/SS36V9_4.1.3/fcm.vmware/t_fvm_ins_ovr.html) • Installing and upgrading IBM Tivoli Storage FlashCopy Manager for Windows (http://www.ibm.com/support/knowledgecenter/SS36V9_4.1.3/fcm.win/c_fcm_installing.html)
Tivoli Storage Manager for Enterprise Resource Planning	• Installing Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP for DB2 (http://www.ibm.com/support/knowledgecenter/SSZHVN_7.1.3/erp.db2/t_dperp_d_install_ovr.html) • Installing Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP for Oracle (http://www.ibm.com/support/knowledgecenter/SSZHVN_7.1.3/erp.orc/t_dperp_o_install_ovr.html)

• To install software on a virtual machine client node, follow the instructions for the selected backup type.

Backup type	Link to instructions
If you plan to create full VMware backups of virtual machines, install and configure the Tivoli Storage Manager backup-archive client.	• Installing the UNIX and Linux backup-archive clients (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/client/c_inst_baunix.html) • Installing the Windows backup-archive client (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/client/t_inst_winclient.html)
If you plan to create incremental forever full backups of virtual machines, install and configure Tivoli Storage Manager for Virtual Environments and a Tivoli Storage Manager backup-archive client on the same client node or on different client nodes.	• Tivoli Storage Manager for Virtual Environments online product documentation (http://www.ibm.com/support/knowledgecenter/SS8TDQ_7.1.3/ve.user/vmware_container_topic.html) **Tip:** You can obtain the software for Tivoli Storage Manager for Virtual Environments and the Tivoli Storage Manager backup-archive client in the Tivoli Storage Manager for Virtual Environments installation package.

2. To allow the client to connect to the server, add or update the values for the **TCPSERVERADDRESS**, **TCPPORT**, and **NODENAME** options in the client options file. Use the values that you recorded when you registered the client ("Registering clients" on page 108).

• For clients that are installed on an AIX, Linux, Mac OS X, or Oracle Solaris operating system, add the values to the client system-options file, dsm.sys.

• For clients that are installed on a Windows operating system, add the values to the dsm.opt file.

By default, the options files are in the installation directory.

3. If you installed a backup-archive client on a Linux or Windows operating system, install the client management service on the client. Follow the instructions in "Installing the client management service" on page 62.

4. Configure the client to run scheduled operations. Follow the instructions in "Configuring the client to run scheduled operations."

5. Optional: Configure communications through a firewall. Follow the instructions in "Configuring client/server communications through a firewall" on page 113.

6. Run a test backup to verify that data is protected as you planned. For example, for a backup-archive client, complete the following steps:

   a. On the Clients page of the Operations Center, select the client that you want to back up, and click **Back Up**.

   b. Verify that the backup completes successfully and that there are no warning or error messages.

7. Monitor the results of the scheduled operations for the client in the Operations Center.

What to do next

If you need to change what is getting backed up from the client, follow the instructions in "Modifying the scope of a client backup" on page 118.

Configuring the client to run scheduled operations

You must configure and start a client scheduler on the client node. The client scheduler enables communication between the client and server so that scheduled operations can occur. For example, scheduled operations typically include backing up files from a client.

About this task

The preferred practice is to install the Tivoli Storage Manager backup-archive client on all client nodes so that you can configure and start the Tivoli Storage Manager client acceptor daemon (CAD) on the client node. The CAD is designed to efficiently run scheduled operations. The CAD manages the client scheduler so that the scheduler runs only when required:

• When it is time to query the server about the next scheduled operation

• When it is time to start the next scheduled operation

By using the CAD, you can reduce the number of background processes on the client and help to avoid memory retention problems.

The CAD runs schedules for the following products: the backup-archive client, Tivoli Storage Manager for Databases, Tivoli Storage Manager for Enterprise Resource Planning, Tivoli Storage Manager for Mail, and Tivoli Storage Manager for Virtual Environments. If you installed a product for which the CAD does not run schedules, follow the configuration instructions in the product documentation to ensure that scheduled operations can occur.

If your business uses a third-party scheduling tool as standard practice, you can use that scheduling tool as an alternative to the CAD. Typically, third-party scheduling tools start Tivoli Storage Manager client programs directly by using operating system commands. To configure a third-party scheduling tool, see the product documentation.

Procedure

To configure and start the client scheduler by using the CAD, follow the instructions for the operating system that is installed on the client node:

AIX and Oracle Solaris

1. From the Tivoli Storage Manager backup-archive client GUI, click **Edit** > **Client Preferences**.
2. Click the **Web Client** tab.
3. In the **Managed Services Options** field, click **Schedule**. If you also want the CAD to manage the web client, click the **Both** option.
4. To ensure that the scheduler can start unattended, in the dsm.sys file, set the **passwordaccess** option to generate.
5. To store the client node password, issue the following command and enter the client node password when prompted:

   `dsmc query sess`
6. Start the CAD by issuing the following command on the command line:

   `/usr/bin/dsmcad`
7. To enable the CAD to start automatically after a system restart, add the following entry to the system startup file (typically, /etc/inittab):

   `tsm::once:/usr/bin/dsmcad > /dev/null 2>&1 # TSM Client Acceptor Daemon`

HP-UX

Follow the instructions in Scheduling events using the command-line client (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/client/t_cfg_cmdlineschedunix.html).

Linux

1. From the Tivoli Storage Manager backup-archive client GUI, click **Edit** > **Client Preferences**.
2. Click the **Web Client** tab.
3. In the **Managed Services Options** field, click **Schedule**. If you also want the CAD to manage the web client, click the **Both** option.
4. To ensure that the scheduler can start unattended, in the dsm.sys file, set the **passwordaccess** option to generate.
5. To store the client node password, issue the following command and enter the client node password when prompted:

   `dsmc query sess`
6. Start the CAD by logging in with the root user ID and issuing the following command:

   `service dsmcad start`
7. To enable the CAD to start automatically after a system restart, add the service by issuing the following command at a shell prompt:

   `# chkconfig --add dsmcad`

MAC OS X

1. In the Tivoli Storage Manager backup-archive client GUI, click **Edit** > **Client Preferences**.
2. To ensure that the scheduler can start unattended, click **Authorization** and select **Password Generate**. Click **Apply**.
3. To specify how services are managed, click **Web Client** and select **Schedule**. Click **Apply** and **OK**.

4. To ensure that the generated password is saved, restart the backup-archive client.
5. Use the Tivoli Storage Manager Tools for Administrators application to start the CAD.

Windows

1. In the Tivoli Storage Manager backup-archive client GUI, click **Utilities** > **Setup Wizard** > **Help me configure the TSM Client Scheduler**. Click **Next**.
2. Read the information on the TSM Scheduler Wizard page and click **Next**.
3. On the Scheduler Task page, select **Install a new or additional scheduler** and click **Next**.
4. On the TSM Scheduler Name and Location page, specify a name for the client scheduler that you are adding. Then, select **Use the Client Acceptor daemon (CAD)** to manage the scheduler and click **Next**.
5. Enter the name that you want to assign to this CAD. The default name is TSM Client Acceptor. Click **Next**.
6. Complete the configuration by stepping through the wizard.
7. Update the client options file, dsm.opt, and set the `passwordaccess` option to `generate`.
8. To store the client node password, issue the following command at the command prompt:

   `dsmc query sess`

   Enter the client node password when prompted.
9. Start the CAD service from the Services Control page. For example, if you used the default name, start the TSM Client Acceptor service. Do not start the scheduler service that you specified on the TSM Scheduler Name and Location page. The scheduler service is started and stopped automatically by the CAD service as needed.

Configuring client/server communications through a firewall

If a client must communicate with a server through a firewall, you must enable client/server communications through the firewall.

About this task

Attention: Do not configure a firewall in a way that might cause termination of sessions that are in use by a server or storage agent. Termination of a valid session can cause unpredictable results. Processes and sessions might appear to stop due to input/output errors. To help exclude sessions from timeout restrictions, configure known ports for Tivoli Storage Manager components. Ensure that the `KEEPALIVE` server option remains set to the default value of YES. In this way, you can help to ensure that client/server communication is uninterrupted. For instructions about setting the `KEEPALIVE` server option, see Technote 1642715 (http://www.ibm.com/support/docview.wss?uid=swg21642715).

Procedure

Open the following ports to allow access through the firewall:

TCP/IP port for the backup-archive client, command-line administrative client, and the client scheduler
Specify the port by using the **tcpport** option in the client options file:

- If you are not using the Secure Sockets Layer (SSL) protocol to secure communications, the **tcpport** option in the client option file must match the **TCPPORT** option in the server options file. The default value is 1500. If you decide to use a value other than the default, specify a number in the range 1024 - 32767.

- If you use the SSL protocol to secure communications, the **tcpport** client option must match the value of the **SSLTCPPORT** server option.

HTTP port to enable communication between the web client and remote workstations
Specify the port for the remote workstation by setting the **httpport** option in the client options file of the remote workstation. The default value is 1581.

TCP/IP ports for the remote workstation
The default value of 0 (zero) causes two free port numbers to be randomly assigned to the remote workstation. If you do not want the port numbers to be randomly assigned, specify values by setting the **webports** option in the client options file of the remote workstation.

TCP/IP port for administrative sessions
Specify the port on which the server waits for requests for administrative client sessions:

- If you are not using the SSL protocol to secure communications, the value of the client **tcpadminport** option must match the value of the **TCPADMINPORT** server option. In this way, you can secure administrative sessions within a private network.

- If you are using the SSL protocol to secure communications, the value of the client **tcpadminport** option must match the value of the server **SSLTCPADMINPORT** option.

Managing client operations

You can evaluate and resolve errors that are related to a backup-archive client by using the Operations Center, which provides suggestions for resolving errors. For errors on other types of clients, you must examine the error logs on the client and review the product documentation.

About this task

In some cases, you can resolve client errors by stopping and starting the client acceptor daemon (CAD). If client nodes or administrator IDs are locked, you can resolve the issue by unlocking the client node or administrator ID, and then resetting the password.

For detailed instructions about identifying and resolving client errors, see Resolving Tivoli Storage Manager client problems (V7.1.1) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.tshoot.doc/r_pdg_clidiagtips.html).

Evaluating errors in client error logs

You can resolve client errors by obtaining suggestions from the Operations Center or by reviewing error logs on the client.

Before you begin

To resolve errors in a backup-archive client on a Linux or Windows operating system, ensure that the client management service is installed and started. For installation instructions, see "Installing the client management service" on page 62. For instructions about verifying the installation, see "Verifying the client management service is installed correctly" on page 62.

Procedure

To diagnose and resolve client errors, take one of the following actions:

- If the client management service is installed on the client node, complete the following steps:
  1. On the Operations Center Overview page, click **Clients** and select the client.
  2. Click **Details**.
  3. On the client Summary page, click the **Diagnosis** tab.
  4. Review the retrieved log messages.

     Tips:
     - To show or hide the Client Logs pane, double-click the Client Logs bar.
     - To resize the Client Logs pane, click and drag the Client Logs bar.

     If suggestions are displayed on the Diagnosis page, select a suggestion. In the Client Logs pane, client log messages to which the suggestion relates are highlighted.
  5. Use the suggestions to resolve the problems that are indicated by the error messages.

     Tip: Suggestions are provided for only a subset of client messages.
- If the client management service is not installed on the client node, review the error logs for the installed client.

Stopping and restarting the client acceptor daemon

If you change the configuration of your solution, you must restart the client acceptor daemon (CAD) on all client nodes where a backup-archive client is installed.

About this task

In some cases, you can resolve client scheduling problems by stopping and restarting the CAD. The CAD must be running to ensure that scheduled operations can occur on the client. For example, if you change the IP address or domain name of the Tivoli Storage Manager server, you must restart the CAD.

Procedure

Follow the instructions for the operating system that is installed on the client node:

AIX and Solaris
- To stop the CAD, complete the following steps:

1. Determine the CAD process ID by issuing the following command on the command line:

   ```
   ps -ef | grep dsmcad
   ```

   Review the output. In the following sample output, 6764 is the CAD process ID:

   ```
   root  6764    1   0 16:26:35 ?          0:00 /usr/bin/dsmcad
   ```

2. Issue the following command on the command line:

   ```
   kill -9 PID
   ```

   where *PID* specifies the CAD process ID.

- To start the CAD, issue the following command on the command line:

  ```
  /usr/bin/dsmcad
  ```

HP-UX

Follow the instructions in Starting and stopping the client service (V7.1.1) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.tshoot.doc/r_pdg_strtstopcli.html).

Linux

- To stop the CAD (and not restart it), issue the following command:

  ```
  # service dsmcad stop
  ```

- To stop and restart the CAD, issue the following command:

  ```
  # service dsmcad restart
  ```

MAC OS X

Click **Applications** > **Utilities** > **Terminal**.

- To stop the CAD, issue the following command:

  ```
  /bin/launchctl unload -w com.ibm.tivoli.dsmcad
  ```

- To start the CAD, issue the following command:

  ```
  /bin/launchctl load -w com.ibm.tivoli.dsmcad
  ```

Windows

- To stop the CAD service, complete the following steps:
  1. Click **Start** > **Administrative Tools** > **Services**.
  2. Double-click the CAD service.
  3. Click **Stop** and **OK**.
- To restart the CAD service, complete the following steps:
  1. Click **Start** > **Administrative Tools** > **Services**.
  2. Double-click the CAD service.
  3. Click **Start** and **OK**.

Related reference:

Resolving client scheduling problems (V7.1.1) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.tshoot.doc/r_pdg_clisched.html)

Resetting passwords

If a password for a client node or an administrator ID is lost or forgotten, you can reset the password. Multiple attempts to access the system with an incorrect password can cause a client node or administrator ID to be locked. You can take steps to resolve the issue.

Procedure

To resolve password issues, take one of the following actions:

- If a backup-archive client is installed on a client node, and the password is lost or forgotten, complete the following steps:

  1. Generate a new password by issuing the **UPDATE NODE** command:

     `update node node_name new_password forcepwreset=yes`

     where *node_name* specifies the client node and *new_password* specifies the password that you assign.

  2. Inform the client node owner about the changed password. When the owner of the client node logs in with the specified password, a new password is generated automatically. That password is unknown to users to enhance security.

     Tip: The password is generated automatically if you previously set the **passwordaccess** option to `generate` in the client options file.

- If an administrator is locked out because of password issues, complete the following steps:

  1. To provide the administrator with access to the server, issue the **UNLOCK ADMIN** command. For instructions, see UNLOCK ADMIN (Unlock an administrator) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_admin_unlock.html).

  2. Set a new password by using the **UPDATE ADMIN** command:

     `update admin admin_name new_password forcepwreset=yes`

     where *admin_name* specifies the name of the administrator and *new_password* specifies the password that you assign.

- If a client node is locked, complete the following steps:

  1. Determine why the client node is locked and whether it must be unlocked. For example, if the client node is decommissioned, the client node is being removed from the production environment. You cannot reverse the decommission operation, and the client node remains locked. A client node also might be locked if the client data is the subject of a legal investigation.

  2. If you must unlock a client node, use the **UNLOCK NODE** command. For instructions, see UNLOCK NODE (Unlock a client node) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_node_unlock.html).

  3. Generate a new password by issuing the **UPDATE NODE** command:

     `update node node_name new_password forcepwreset=yes`

     where *node_name* specifies the name of the node and *new_password* specifies the password that you assign.

4. Inform the client node owner about the changed password. When the owner of the client node logs in with the specified password, a new password is generated automatically. That password is unknown to users to enhance security.

   Tip: The password is generated automatically if you previously set the `passwordaccess` option to `generate` in the client options file.

Modifying the scope of a client backup

When you set up client backup operations, the preferred practice is to exclude objects that you do not require. For example, you typically want to exclude temporary files from a backup operation.

About this task

When you exclude unnecessary objects from backup operations, you get better control of the amount of storage space that is required for backup operations, and the cost of storage. Depending on your licensing package, you also might be able to limit licensing costs.

Procedure

How you modify the scope of backup operations depends on the product that is installed on the client node:

- For a backup-archive client, you can create an include-exclude list to include or exclude a file, groups of files, or directories from backup operations. To create an include-exclude list, follow the instructions in Creating an include-exclude list (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/client/t_cfg_crtinclexcl.html).

  To ensure consistent use of an include-exclude list for all clients of one type, you can create a client option set on the server that contains the required options. Then, you assign the client option set to each of the clients of the same type. For details, see Controlling client operations through client option sets (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.admin/t_define_client_optsets.html).

- For a backup-archive client, you can specify the objects to include in an incremental backup operation by using the `domain` option. Follow the instructions in Domain client option (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/client/r_opt_domain.html).

- For other products, to define which objects are included in and excluded from backup operations, follow the instructions in the product documentation.

Managing client upgrades

When a fix pack or interim fix becomes available for a client, you can upgrade the client to take advantage of product improvements. Servers and clients can be upgraded at different times and can be at different levels with some restrictions.

Before you begin

1. Review the client/server compatibility requirements in Technote 1053218 (http://www.ibm.com/support/docview.wss?uid=swg21053218). If your solution includes servers or clients at a level that is earlier than V7.1, review the guidelines to ensure that client backup and archive operations are not disrupted.

2. Verify system requirements for the client in Tivoli Storage Manager Supported Operating Systems (http://www.ibm.com/support/docview.wss?uid=swg21243309).

3. If the solution includes storage agents or library clients, review the information about storage-agent and library-client compatibility with Tivoli Storage Manager servers that are configured as library managers. See Technote 1302789 (http://www.ibm.com/support/docview.wss?uid=swg21302789).

   If you plan to upgrade a library manager and a library client, you must upgrade the library manager first.

Procedure

To upgrade the software, follow the instructions.

Software	Link to instructions
Tivoli Storage Manager backup-archive client	• Upgrading the backup-archive client (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/client/c_inst_mig.html)
Tivoli Storage FlashCopy Manager	• Installing and upgrading IBM Tivoli Storage FlashCopy Manager for UNIX and Linux (http://www.ibm.com/support/knowledgecenter/SS36V9_4.1.3/fcm.unx/t_protect_fcmul.html) • Installing and upgrading Tivoli Storage FlashCopy Manager for VMware (http://www.ibm.com/support/knowledgecenter/SS36V9_4.1.3/fcm.vmware/t_fvm_ins_ovr.html) • Installing and upgrading IBM Tivoli Storage FlashCopy Manager for Windows (http://www.ibm.com/support/knowledgecenter/SS36V9_4.1.3/fcm.win/c_fcm_installing.html)
Tivoli Storage Manager for Databases	• Upgrading Data Protection for SQL Server (http://www.ibm.com/support/knowledgecenter/SSTFZR_7.1.3/db.sql/t_dpsql_upgrade.html) • Data Protection for Oracle installation (http://www.ibm.com/support/knowledgecenter/SSTFZR_7.1.3/db.orc/t_dporc_installing.html) • Installing, upgrading, and migrating IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server (http://www.ibm.com/support/knowledgecenter/SSTG2D_7.1.3/mail.exc/t_dpexc_install_ovr.html)
Tivoli Storage Manager for Enterprise Resource Planning	• Upgrading Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP for DB2 (http://www.ibm.com/support/knowledgecenter/SSZHVN_7.1.3/erp.db2/t_dperp_d_install_upgrade.html) • Upgrading Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP for Oracle (http://www.ibm.com/support/knowledgecenter/SSZHVN_7.1.3/erp.orc/t_dperp_o_install_upgrade.html)
Tivoli Storage Manager for Mail	• Installation of Data Protection for IBM Domino on a UNIX, AIX, or Linux system (http://www.ibm.com/support/knowledgecenter/SSTG2D_7.1.0/com.ibm.itsm.mail.dom.unx.doc/concept_dominstall.html) • Installation of Data Protection for IBM Domino on a Windows system (http://www.ibm.com/support/knowledgecenter/SSTG2D_7.1.0/com.ibm.itsm.mail.dom.win.doc/concept_dominstall.html) • Installing, upgrading, and migrating IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server (http://www.ibm.com/support/knowledgecenter/SSTG2D_7.1.3/mail.exc/t_dpexc_install_ovr.html)

Software	Link to instructions
Tivoli Storage Manager for Virtual Environments	• Installing and upgrading Data Protection for VMware (http://www.ibm.com/support/knowledgecenter/SS8TDQ_7.1.3/ve.inst/t_ve_installing.html) • Installing Data Protection for Microsoft Hyper-V (http://www.ibm.com/support/knowledgecenter/SS8TDQ_7.1.3/ve.hv/c_ve_hv_inst_plan.html)

Decommissioning a client node

If a client node is no longer required, you can start a process to remove it from the production environment. For example, if a workstation was backing up data to a Tivoli Storage Manager server, but the workstation is no longer used, you can decommission the workstation.

About this task

When you start the decommission process, the server locks the client node to prevent it from accessing the server. Files that belong to the client node are gradually deleted, and then the client node is deleted. You can decommission the following types of client nodes:

Application client nodes
> Application client nodes include email servers, databases, and other applications. For example, any of the following applications can be an application client node:
> - FlashCopy Manager
> - Tivoli Storage Manager for Databases
> - Tivoli Storage Manager for Enterprise Resource Planning
> - Tivoli Storage Manager for Mail
> - Tivoli Storage Manager for Virtual Environments

System client nodes
> System client nodes include workstations, network-attached storage (NAS) file servers, and API clients.

Virtual machine client nodes
> Virtual machine client nodes consist of an individual guest host within a hypervisor. Each virtual machine is represented as a Tivoli Storage Manager file space.

The simplest method for decommissioning a client node is to use the Operations Center. The decommission process runs in the background. If the client is configured to replicate client data, the Operations Center automatically removes the client from replication on the source and target replication servers before it decommissions the client.

Tip: Alternatively, you can decommission a client node by issuing the `DECOMMISSION NODE` or `DECOMMISSION VM` command. You might want to use this method in the following cases:
- To schedule the decommission process for the future or to run a series of commands by using a script, specify the decommission process to run in the background.

- To monitor the decommission process for debugging purposes, specify the decommission process to run in the foreground. If you run the process in the foreground, you must wait for the process to be completed before you continue with other tasks.

Procedure

Take one of the following actions:
- To decommission a client in the background by using the Operations Center, complete the following steps:
  1. On the Operations Center Overview page, click **Clients** and select the client.
  2. Click **More** > **Decommission**.
- To decommission a client node by using an administrative command, complete the following steps:
  1. Determine whether the client node is configured for node replication by issuing the **QUERY NODE** command. For example, if the client node is named AUSTIN, run the following command:

     ```
     query node austin format=detailed
     ```

     Review the Replication State output field.
  2. If the client node is configured for replication, remove the client node from replication by issuing the **REMOVE REPLNODE** command. For example, if the client node is named AUSTIN, issue the following command:

     ```
     remove replnode austin
     ```
  3. Take one of the following actions:
     - To decommission an application or system client node in the background, issue the **DECOMMISSION NODE** command. For example, if the client node is named AUSTIN, issue the following command:

       ```
       decommission node austin
       ```
     - To decommission an application or system client node in the foreground, issue the **DECOMMISSION NODE** command and specify the `wait=yes` parameter. For example, if the client node is named AUSTIN, issue the following command:

       ```
       decommission node austin wait=yes
       ```
     - To decommission a virtual machine in the background, issue the **DECOMMISSION VM** command. For example, if the virtual machine is named AUSTIN, the file space is 7, and the file space name is specified by the file space ID, issue the following command:

       ```
       decommission vm austin 7 nametype=fsid
       ```
     - To decommission a virtual machine in the foreground, issue the **DECOMMISSION VM** command and specify the `wait=yes` parameter. For example, issue the following command:

       ```
       decommission vm austin 7 nametype=fsid wait=yes
       ```

Related reference:

➡ DECOMMISSION NODE (Decommission a client node) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_node_decommission.html)

➡ DECOMMISSION VM (Decommission a virtual machine) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_vm_decommission.html)

QUERY NODE (Query nodes) (http://www.ibm.com/support/
knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_node_query.html)

REMOVE REPLNODE (Remove a client node from replication)
(http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/
r_cmd_replnode_remove.html)

Deactivating data to free storage space

In some cases, you can deactivate data that is stored on a Tivoli Storage Manager
server. When you run the deactivation process, any backup data that was stored
before the specified date and time is deactivated and will be deleted as it expires.
In this way, you can free space on the server.

About this task

Some application clients always save data to the server as active backup data.
Because active backup data is not managed by inventory expiration policies, the
data is not deleted automatically, and uses server storage space indefinitely. To free
the storage space that is used by obsolete data, you can deactivate the data.

When you run the deactivation process, all active backup data that was stored
before the specified date becomes inactive. The data is deleted as it expires and
cannot be restored. The deactivation feature applies only to application clients that
protect Oracle databases.

Procedure

1. From the Operations Center Overview page, click **Clients**.
2. In the Clients table, select one or more clients and click **More** > **Clean Up**.

   Command-line method: Deactivate data by using the `DEACTIVATE DATA`
   command.

Related reference:

DEACTIVATE DATA (Deactivate data for a client node) (http://
www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/
r_cmd_data_deactivate.html)

Chapter 19. Managing data storage

Managing your data for efficiency and adding to Tivoli Storage Manager supported devices and media to store your data.

Related reference:

Comparing storage pools (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.admin/r_stgpools_compare.html)

Auditing a storage pool

Audit a storage pool container to check for inconsistencies between database information and a container.

About this task

You audit a storage pool container in the following situations:
- When you issue the **QUERY DAMAGED** command and a problem is detected.
- When the server displays messages about damaged data extents.
- Your hardware reports an issue and error messages that are associated with the storage pool container are displayed.

Procedure

To audit a storage pool container, issue the **AUDIT CONTAINER** command. For example, issue the following command to audit a container, `000000000000076c.dcf`:

```
audit container c:\tsm-storage\07\000000000000076c.dcf
```

The output from the ANR4891I message displays information about any damaged data extents.

What to do next

If you detect problems with the storage pool container, you can restore data based on your configuration. You can repair the contents in the storage pool by using the **REPAIR STGPOOL** command.

Restriction: You can only repair the contents of the storage pool if you protected the storage pool by using the **PROTECT STGPOOL** command.

Related reference:

AUDIT CONTAINER (Verify the consistency of database information for a directory-container storage pool) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_container_audit.html)

QUERY DAMAGED (Query damaged data in a directory-container or cloud-container storage pool) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_damaged_query.html)

Managing inventory capacity

Manage the capacity of the database, active log, and archive logs to ensure that the inventory is sized for the tasks, based on the status of the logs.

Before you begin

The active and archive logs have the following characteristics:

- The active log can be a maximum size of 512 GB. For more information about the sizing the active log for your system, see Planning the storage arrays (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.solutions/r_msdisk_arrays_plan.html)
- The archive log size is limited to the size of the file system that it is installed on. The archive log size is not maintained at a predefined size like the active log. Archive log files are automatically deleted after they are no longer needed.

As a best practice, you can optionally create an archive failover log to store archive log files when the archive log directory is full.

Check the Operations Center to determine the component of the inventory that is full. Ensure that you stop the server before you increase the size of one of the inventory components.

Procedure

1. To increase the size of the database, complete the following steps:
   - Create one or more directories for the database on separate drives or file systems.
   - Issue the **EXTEND DBSPACE** command to add the directory or directories to the database. The directories must be accessible to the instance user ID of the database manager. By default, data is redistributed across all database directories and space is reclaimed.

     Tips:
     - The time that is needed to complete redistribution of data and reclaiming of space is variable, depending on the size of your database. Make sure that you plan adequately.
     - Ensure that the directories that you specify are the same size as existing directories to ensure a consistent degree of parallelism for database operations. If one or more directories for the database are smaller than the others, they reduce the potential for optimized parallel prefetching and distribution of the database.
   - Halt and restart the server to fully use the new directories.
   - Reorganize the database if necessary. Index and table reorganization for the server database can help to avoid unexpected database growth and performance issues. For more information about reorganizing the database, see Technote 1683633 (http://www.ibm.com/support/docview.wss?uid=swg21683633).

2. To decrease the size of the database for V7.1 servers and later, issue the following DB2 commands from the server instance directory:

   Restriction: The commands can increase I/O activity, and might affect server performance. To minimize performance problems, wait until one command is completed before you issue the next command. The DB2 commands can be issued when the server is running.

```
db2 connect to tsmdb1
db2 set schema tsmdb1
db2 ALTER TABLESPACE USERSPACE1 REDUCE MAX
db2 ALTER TABLESPACE IDXSPACE1 REDUCE MAX
db2 ALTER TABLESPACE LARGEIDXSPACE1 REDUCE MAX
db2 ALTER TABLESPACE LARGESPACE1 REDUCE MAX
db2 ALTER TABLESPACE REPLTBLSPACE1 REDUCE MAX
db2 ALTER TABLESPACE REPLIDXSPACE1 REDUCE MAX
db2 ALTER TABLESPACE ARCHOBJDATASPACE REDUCE MAX
db2 ALTER TABLESPACE ARCHOBJIDXSPACE REDUCE MAX
db2 ALTER TABLESPACE BACKOBJDATASPACE REDUCE MAX
db2 ALTER TABLESPACE BACKOBJIDXSPACE REDUCE MAX
db2 ALTER TABLESPACE BFABFDATASPACE REDUCE MAX
db2 ALTER TABLESPACE BFABFIDXSPACE REDUCE MAX
db2 ALTER TABLESPACE BFBFEXTDATASPACE REDUCE MAX
db2 ALTER TABLESPACE BFBFEXTIDXSPACE REDUCE MAX
db2 ALTER TABLESPACE DEDUPTBLSPACE1 REDUCE MAX
db2 ALTER TABLESPACE DEDUPIDXSPACE1 REDUCE MAX
db2 ALTER TABLESPACE DEDUPTBLSPACE2 REDUCE MAX
db2 ALTER TABLESPACE DEDUPIDXSPACE2 REDUCE MAX
db2 ALTER TABLESPACE DEDUPTBLSPACE3 REDUCE MAX
db2 ALTER TABLESPACE DEDUPIDXSPACE3 REDUCE MAX
db2 ALTER TABLESPACE DEDUPTBLSPACE4 REDUCE MAX
db2 ALTER TABLESPACE DEDUPIDXSPACE4 REDUCE MAX
db2 ALTER TABLESPACE DEDUPTBLSPACE5 REDUCE MAX
db2 ALTER TABLESPACE DEDUPIDXSPACE5 REDUCE MAX
```

3. To increase or decrease the size of the active log, complete the following steps:

   a. Ensure that the location for the active log has enough space for the increased log size. If a log mirror exists, its location must also have enough space for the increased log size.

   b. Halt the server.

   c. In the `dsmserv.opt` file, update the **ACTIVELOGSIZE** option to the new size of the active log, in megabytes.

      The size of an active log file that is used by Tivoli Storage Manager is based on the value of the `ACTIVELOGSIZE` option. Guidelines for space requirements are in the following table:

Table 18. How to estimate volume and file space requirements

ACTIVELOGSize option value	**Reserve this much free space in the active log directory, in addition to the ACTIVELOGSize space**
2 GB - 128 GB	5120 MB
129 GB - 256 GB	10240 MB
257 GB - 512 GB	20480 MB

   To change the active log to its maximum size of 512 GB, enter the following server option:

   ```
   activelogsize 524288
   ```

   d. If you plan to use a new active log directory, update the directory name that is specified in the **ACTIVELOGDIRECTORY** server option. The new directory must be empty and must be accessible to the user ID of the database manager.

   e. Restart the server.

4. Compress the archive logs to reduce the amount of space that is required for storage. Enable dynamic compression of the archive log by issuing the following command:

   ```
   SETOPT  ARCHLOGCOMPress Yes
   ```

Restriction: Use caution when you enable the `ARCHLOGCOMPRESS` server option on systems with sustained high volume usage and heavy workloads. Enabling this option in this system environment can cause delays in archiving log files from the active log file system to the archive log file system. This delay can cause the active log file system to run out of space. Be sure to monitor the available space in the active log file system after archive log compression is enabled. If the active log directory file system usage nears out of space conditions, the `ARCHLOGCOMPRESS` server option must be disabled. You can use the `SETOPT` command to disable archive log compression immediately without halting the server.

Related reference:

➡ ACTIVELOGSIZE (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_opt_server_activelogsize.html)

➡ EXTEND DBSPACE (Increase space for the database) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_dbspace_extend.html)

➡ SETOPT (Set a server option for dynamic update) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_setopt.html)

Managing memory and CPU usage

Ensure that you manage memory requirements and Central Processing Unit (CPU) usage to ensure that Tivoli Storage Manager server can complete data processes such as backup and data deduplication. Consider the impact on performance when you complete certain processes.

Before you begin

- Ensure that your configuration uses the minimum hardware and software requirements. For more information, see Tivoli Storage Manager Supported Operating Systems (http://www.ibm.com/support/docview.wss?uid=swg21243309).
- For more information about managing resources such as the database and recovery log, see Planning the storage arrays (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.solutions/r_msdisk_arrays_plan.html).
- Add more system memory to determine whether there is a performance improvement. Monitor memory usage regularly to determine whether more memory is required.

Procedure

1. Release memory from the file system cache where possible.
2. To manage the system memory that is used by each server on a system, use the `DBMEMPERCENT` server option. Limit the percentage of system memory that can be used by the database manager of each server. If all servers are equally important, use the same value for each server. If one server is a production server and other servers are test servers, set the value for the production server to a higher value than the test servers.
3. Set the user data limit and private memory for the database to ensure that private memory is not exhausted. Exhausting private memory can result in errors, less than optimal performance, and instability.

Tuning scheduled activities

Schedule maintenance tasks daily to ensure that your solution operates correctly. By tuning your solution, you maximize Tivoli Storage Manager server resources and effectively use different functions available within your solution.

Procedure

1. Monitor system performance regularly to ensure that client backup and server maintenance tasks are completing successfully. Follow the instructions in Part 3, "Monitoring a multisite disk solution," on page 71.

2. Optional: If monitoring information shows that the workload that the server is handling has increased, review the planning information. An increase in the number of clients, the amount of data being backed up, or the amount of time that is available for backups can mean that you must modify the capacity of the system is adequate.

3. Determine whether your solution is performing at the level you expect. Review the client schedule to check whether tasks are completing within the scheduled time frame:

   a. On the **Clients** page of the Operations Center, select the client.

   b. Click **Details**.

   c. From the client Summary page, review the **Backup Up** and **Replication** activity to identify any risks.

   Adjust the time and frequency of client backup operations, if necessary.

4. Schedule enough time for the following maintenance tasks to complete successfully within a 24-hour period:

   a. Protect storage pools.

   b. Replicate node data.

   c. Back up the database.

   d. Run expiration processing to remove client backups and archive file copies from server storage.

   Tip: Schedule maintenance tasks to start at an appropriate time and in the correct sequence. For example, schedule replication tasks after client backups complete successfully.

Related concepts:

⬛➡ Performance (V7.1.1) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.perf.doc/c_performance.html)

Related tasks:

⬛➡ Deduplicating data (V7.1.1) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.srv.doc/t_dedup.html)

"Defining schedules for server maintenance activities" on page 58

Moving clients from one server to another

To avoid running out of space on a server or to resolve workload issues, you might have to move client nodes from one Tivoli Storage Manager server to another.

Before you begin

Plan the capacity for your solution to ensure that you have enough space for client nodes on the server, which includes space for future growth.

About this task

When you move the client nodes, you can leave their existing backups on the original server to expire according to your expiration policy, or export their existing backups to the new server.

Procedure

Complete the following steps to move a client node to another server.

1. Export the client node directly to a new server by using the **EXPORT NODE** command.
2. Update the client options file with the new server name.
3. On the new server, assign a schedule for the client node to back up data.
   a. On the Operations Center **Clients** page, select the client node.
   b. Click **More** > **Schedule Association**.
   c. Select the check box in the rows of schedule to which you want to assign the selected client node.
   d. Click **Save**.
4. Issue the **EXPORT NODE** command again to incrementally export data from the original server to the new server. By incrementally exporting data, you export data that was backed up between the first export process and when you assigned a schedule to the client node.
5. Monitor the client node to ensure that it is backing up data according to the schedule that you set and to monitor whether the client node is at risk. Hover over **Clients** and click **Schedules**.
6. Decommission the client node from the original server by completing the following steps.
   a. On the Operations Center **Overview** page, click **Clients**.
   b. In the Clients table, select the client node.
   c. Click **More** > **Decommission**.

   The client node is removed from the original server. As the data expires, as specified in your policy settings, the client node data is deleted. After the client node data is deleted, the client is removed from the server.

Related reference:

EXPORT NODE (Export client node information) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_node_export.html)

IMPORT NODE (Import client node information) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_node_import.html)

Chapter 20. Managing replication

Use replication to recover data at a disaster recovery site and to maintain the same level of files on the source and target servers. You can manage replication at node level and you can also protect data at storage-pool level.

Enabling node replication

With Tivoli Storage Manager, you can enable node replication to protect your data.

About this task

Replicate the client node if you want to replicate all client data, including metadata. By default, node replication is disabled when you start the Tivoli Storage Manager server for the first time.

Tips:
- As a best practice, protect the storage pool before you replicate the client node. When node replication is started, the data extents that are already replicated through storage pool protection are skipped, which reduces replication processing time.
- Replication requires increased amounts of memory and there must be sufficient bandwidth to complete processing. Size the database and its logs to ensure that transactions can complete.

Procedure

To enable node replication, complete the following steps in the Operations Center:
1. On the Operations Center Servers page, click **Details**.
2. On the Details page, click **Properties**.
3. In the **Replication** section, select **Enabled** in the **Outbound replication** field.
4. Click **Save**.

What to do next

To verify that replication was successful, see Chapter 13, "Daily monitoring checklist," on page 73.

Protecting data in storage pools

Use storage pool protection to protect your data in directory-container storage pools.

About this task

By protecting a storage pool, you do not use resources that replicate existing data and metadata, which improves server performance. You must use directory-container storage pools if you want to protect and back up the storage pool only.

As a best practice, protect the storage pool before you replicate the client node. When node replication is started, the data extents that are already replicated through storage pool protection are skipped, which reduces the replication processing time.

Procedure

To enable storage pool replication, issue the **PROTECT STGPOOL** command from the source server to back up data extents in a directory-container storage pool. For example, to protect a directory-container storage pool that is named POOL1 issue the following command:

```
protect stgpool pool1
```

Restriction: You can repair the storage pool, by using the **REPAIR STGPOOL** command, if the storage pool is protected. If you replicate the node and you do not protect the storage pool, you cannot repair the storage pool.

What to do next

To verify that replication was successful, see Chapter 13, "Daily monitoring checklist," on page 73.

Modifying replication settings

Modify replication settings in the Operations Center. Change settings such as the number of replication sessions, replication rules, the data that you want to replicate, the replication schedule, and the replication workload.

About this task

You might need to customize your replication settings in the following scenarios:
- Changes to data priorities
- Scheduled processes that negatively affect server performance
- Requirement for a different or extra server to be the target server
- Changes to data retention policies

Procedure

Complete the following tasks to modify replication settings.

Table 19. Replication tasks and procedures

Task	Procedure
Change a replication rule.	1. On the Operations Center Servers page, click **Details**. 2. On the Details page, click **Properties**. 3. In the **Replication** section, choose the replication rule that you want to apply from the **Default archive rule**, **Default backup rule**, and **Default space-management rule** drop-down menu. 4. Click **Save**. To set different replication policies, follow the instructions in "Setting different retention policies for the source server and target server" on page 132.
Specify the duration that replication records are retained.	1. On the Operations Center Servers page, click **Details**. 2. On the Details page, click **Properties**. 3. In the **Replication** section, enter the number of days that replication records must be retained in the **Retain replication history** field. Alternatively, select the **Do not retain** check box if you do not require replication records. 4. Click **Save**.
Specify a target replication server.	1. On the Operations Center Servers page, click **Details**. 2. On the Details page, click **Properties**. 3. In the **Replication** section, specify the target server. 4. Click **Save**.
Cancel replication.	1. On the Operations Center Servers page, click **Active tasks**. 2. Select the process or session that you want to cancel. 3. Click **Cancel**.

Setting different retention policies for the source server and target server

You can maintain different versions of files on the source replication server and the target replication server. To use dissimilar policies, change replication policies on the target replication server independently from the source replication server.

Procedure

1. On the Operations Center Servers page, click **Details**.
2. On the Details page, click **Properties**.
3. In the Replication section, choose the replication rule from the following drop-down menu that you want to apply to the source replication server:
   - Default archive rule
   - Default backup rule
   - Default space-management rule
4. Click **Save**.

Chapter 21. Securing the Tivoli Storage Manager server

Secure the Tivoli Storage Manager server and data by controlling access to servers and client nodes, encrypting data, and maintaining secure access levels and passwords.

Security concepts

You can protect Tivoli Storage Manager from security risks by using communication protocols, securing passwords, and providing different access levels for administrators.

Transport Layer Security

You can use the Secure Sockets Layer (SSL) or the Transport Layer Security (TLS) protocol to provide transport layer security for a secure connection between Tivoli Storage Manager servers, clients, and storage agents. If you send data between the server, client, and storage agent, use SSL or TLS to encrypt the data.

Tip: Any Tivoli Storage Manager documentation that indicates "SSL" or to "select SSL" applies to TLS.

SSL is provided by the Global Security Kit (GSKit) that is installed with the Tivoli Storage Manager server that the server, client, and storage agent use.

Restriction: Do not use the SSL or TLS protocols for communications with a DB2 database instance that is used by a Tivoli Storage Manager server.

Each Tivoli Storage Manager server, client, or storage agent that enables SSL must use a trusted self-signed certificate or obtain a unique certificate that is signed by a certificate authority (CA). You can use your own certificates or purchase certificates from a CA. Either certificate must be installed and added to the key database on the Tivoli Storage Manager server, client, or storage agent. The certificate is verified by the SSL client or server that requests or initiates the SSL communication. Some CA certificates are preinstalled in the key databases, by default.

SSL is set up independently on the Tivoli Storage Manager server, client, and storage agent.

Tivoli Storage Manager authority levels

With Tivoli Storage Manager server, there are different administrative authority levels available that determine the different tasks that an administrator can complete.

After registration, an administrator must be granted authority by being assigned one or more administrative authority levels. An administrator with system authority can complete any task with the server and assign authority levels to other administrators by using the **GRANT AUTHORITY** command. Administrators with policy, storage, or operator authority can complete subsets of tasks.

An administrator can register other administrator IDs, grant levels of authority to them, rename IDs, remove IDs, and lock and unlock them from the Tivoli Storage Manager server.

An administrator can control access to specific client nodes for root user IDs and non-root user IDs. By default a non-root user ID cannot back up data on the node. Use the **UPDATE NODE** command to change the node settings to enable backup.

Tivoli Storage Manager passwords

By default, the server automatically sets password authentication to "on." With password authentication set to on, all users must enter a password when you access the Tivoli Storage Manager server.

Use Lightweight Directory Access Protocol (LDAP) to apply stricter requirements for passwords. For more information, see Managing passwords and logon procedures (V7.1.1) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.srv.doc/c_mgclinod_managepwlogin.html).

Table 20. Password authentication characteristics

Characteristic	More information
Case-sensitivity	Not case-sensitive.
Default password expiration	90 days. The expiration period begins when an administrator ID or client node is first registered to the server. If a user's password is not changed within this period, the user must change the password the next time they access the server.
Invalid password attempts	You can set a limit on consecutive invalid password attempts for all client nodes. When the limit is exceeded, the server locks the node.
Password length	The administrator can specify a minimum password length that is required for Tivoli Storage Manager passwords.

Related reference:

Securing communications (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.admin/r_comms_securing.html)

Managing administrators

An administrator with system authority can complete any task with the server, including assigning authority levels to other administrators. To complete some Tivoli Storage Manager tasks, you must be granted authority by being assigned one or more authority levels.

Procedure

Complete the following tasks to modify administrator settings.

Task	Procedure

Add an administrator.	To add an administrator, ADMIN1, with system authority and specify a password, complete the following steps:
	
	1. Register the administrator and specify **Pa$#$twO** as the password by issuing the following command:
	
	`register admin admin1 Pa$#$twO`
	
	2. Grant system authority to the administrator by issuing the following command:
	
	`grant authority admin1 classes=system`

Change administrative authority.	Change the authority level for an administrator, ADMIN1. • Grant system authority to the administrator by issuing the following command: `grant authority admin1 classes=system` • Revoke authority for the administrator by issuing the following command:
Remove administrators.	Remove an administrator, ADMIN1, from accessing the Tivoli Storage Manager server by issuing the following command: `remove admin admin1`
Temporarily prevent access to the server.	Lock or unlock an administrator by using the **LOCK ADMIN** or **UNLOCK ADMIN** command.

Changing password requirements

You can change the minimum password limit, password length, password expiration, and enable or disable authentication.

About this task

By enforcing password authentication and managing password restrictions, you protect your data and your servers from potential security risks.

Procedure

Complete the following tasks to change password requirements for Tivoli Storage Manager servers.

Table 21. Authentication tasks for Tivoli Storage Manager servers

Task	Procedure for Tivoli Storage Manager server
Set a limit for invalid password attempts.	1. On the TSM Servers page in the Operations Center, select the server whose properties you want to change. 2. Click **Details**, and then click the **Properties** tab. 3. Set the number of invalid attempts in the **Invalid sign-on attempt limit** field. The default value at installation is 0.
Set a minimum password limit.	1. On the TSM Servers page in the Operations Center, select the server whose properties you want to change. 2. Click **Details** and then click the **Properties** tab. 3. Set the number of number of characters in the **Minimum password length** field.
Set password expiration limit.	1. On the TSM Servers page in the Operations Center, select the server whose properties you want to change. 2. Click **Details** and then click the **Properties** tab. 3. Set the number of days in the **Password common expiration** field.
Disable password authentication.	By default, the server automatically sets password authentication to "on." With password authentication set to on, all users must enter a password to access the server. You can disable password authentication only for passwords that authenticate with the Tivoli Storage Manager server (LOCAL). By disabling password authentication, you expose the Tivoli Storage Manager server to security risks.

Set a default authentication method.	Issue the **SET DEFAULTAUTH** command:
	`set defaultauth=local`
	
	To update a client node that authenticates with the Tivoli Storage Manager server, include AUTHENTICATION=LOCAL in the **UPDATE NODE** command:
	`update node authentication=local`

Related concepts:

➡ Managing passwords and logon procedures (V7.1.1) (http://www.ibm.com/
support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.srv.doc/
c_mgclinod_managepwlogin.html)

Securing Tivoli Storage Manager on the system

Protect the system where the Tivoli Storage Manager server runs to prevent unauthorized access.

Procedure

Ensure that unauthorized users cannot access the database directory and the server instance directory. Keep the access settings for these directories that you configured during implementation.

Restricting user access to the server

Authority levels determine what an administrator can do with the Tivoli Storage Manager server. An administrator with system authority can complete any task with the server. Administrators with policy, storage, or operator authority can complete subsets of tasks.

Procedure

1. After you register an administrator by using the **REGISTER ADMIN** command, use the **GRANT AUTHORITY** command to set the administrator's authority level. For details about setting and changing authority, see "Managing administrators" on page 134.
2. To control the authority of an administrator to complete some tasks, use the following two server options:
   a. You can select the authority level that an administrator must have to issue **QUERY** and **SELECT** commands with the **QUERYAUTH** server option. By default, no authority level is required. You can change the requirement to one of the authority levels, including system.
   b. You can specify that system authority is required for commands that cause the server to write to an external file with the **REQSYSAUTHOUTFILE** server option. By default, system authority is required for such commands.
3. You can restrict data backup on a client node to only root user IDs or authorized users. For example, to limit backups to the root user ID, issue the **REGISTER NODE** or **UPDATE NODE** command and specify the **BACKUPINITIATION=root** parameter:

   ```
   update node backupinitiation=root
   ```

Limiting access through port restrictions

Limit access to Tivoli Storage Manager by applying port restrictions.

About this task

You might need to restrict access to specific Tivoli Storage Manager servers, based on your business and security requirements. The Tivoli Storage Manager server can be configured to listen on four TCP/IP ports: two for regular protocols and two for the Transport Layer Security (TLS) protocols.

Procedure

You can set the following server options to specify the port that you require:

Table 22. Server options and port access

Server option	Port access
SSLTCPPORT	Specifies the SSL TCP/IP port address for a Tivoli Storage Manager server. There is no default port value. If the client is configured for SSL communication, the client uses the SSL port to communicate with the target server during failover.
SSLTCPADMINPORT	Specifies the port address on which the server TCP/IP communication driver waits for requests for SSL-enabled sessions. There is no default port value. Use this option to separate administrative client traffic from regular client traffic that uses the TCPPORT and SSLTCPPORT options.
TCPPORT	Specifies the port number on which the server TCP/IP communication driver is to wait for requests for client sessions. The default value is 1500.
TCPADMINPORT	Specifies the port number on which the server TCP/IP communication driver is to wait for requests for sessions other than client sessions. The default is the value of TCPPORT. Use this option to separate administrative client traffic from regular client traffic that uses the TCPPORT and SSLTCPPORT options.

Chapter 22. Stopping and starting the server

Before you complete maintenance or reconfiguration tasks, stop the Tivoli Storage Manager server. Then, start the server in maintenance mode. When you are finished with the maintenance or reconfiguration tasks, restart the server in production mode.

Before you begin

You must have system or operator privilege to stop and start the server.

Stopping the server

Before you stop the server, prepare the system by ensuring that all database backup operations are completed, and all other processes and sessions are ended. In this way, you can safely shut down the server and ensure that data is protected.

About this task

When you issue the **HALT** command to stop the server, the following actions occur:
- All processes and client node sessions are canceled.
- All current transactions are stopped. (The transactions will be rolled back when the server is restarted.)

Procedure

To prepare the system and stop the server, complete the following steps:
1. Prevent new client node sessions from starting by issuing the **DISABLE SESSIONS** command:

   ```
   disable sessions all
   ```
2. Determine whether any client node sessions are in progress by issuing the **QUERY SESSION** command:

   ```
   query session
   ```

   If any client node sessions are in progress, a table displays the session numbers.
3. Wait until the client node sessions are completed or cancel them:
   - To cancel all current sessions, issue the **CANCEL SESSION** command:

     ```
     cancel session all
     ```
   - To cancel a specific session, issue the following command:

     ```
     cancel session session_number
     ```

     where *session_number* specifies the session number.
4. Determine whether any processes are in progress by issuing the **QUERY PROCESS** command:

   ```
   query process
   ```

   If a database backup process is in progress, wait until the process is completed.
5. If other processes are in progress, wait until they are completed or cancel them by issuing the **CANCEL PROCESS** command:

```
cancel process process_number
```

where *process_number* specifies the number of the process.

Tip: If the process that you want to cancel is waiting for a tape volume to be mounted, the mount request is canceled. For example, if you issue an **EXPORT**, **IMPORT**, or **MOVE DATA** command, the command might initiate a process that requires a tape volume to be mounted. However, if a tape volume is being mounted by an automated library, the cancel operation might not take effect until the mount process is complete. Depending on your system environment, this could take several minutes.

6. Stop the server by issuing the **HALT** command:

```
halt
```

Starting the server for maintenance or reconfiguration tasks

Before you begin server maintenance or reconfiguration tasks, start the server in maintenance mode. When you start the server in maintenance mode, you disable operations that might disrupt your maintenance or reconfiguration tasks.

About this task

Start the server in maintenance mode by running the **DSMSERV** utility with the **MAINTENANCE** parameter.

The following operations are disabled in maintenance mode:
- Administrative command schedules
- Client schedules
- Reclamation of storage space on the server
- Inventory expiration
- Migration of storage pools

In addition, clients are prevented from starting sessions with the server.

Tips:
- You do not have to edit the server options file, dsmserv.opt, to start the server in maintenance mode.
- While the server is running in maintenance mode, you can manually start the storage-space reclamation, inventory expiration, and storage-pool migration processes.

Procedure

To start the server in maintenance mode, issue the following command:

```
dsmserv maintenance
```

What to do next

To resume server operations in production mode, complete the following steps:

1. Shut down the server by issuing the **HALT** command:

```
halt
```

2. Start the server by using the method that you use in production mode. Follow the instructions for your operating system:

- ` AIX ` Starting the server instance (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.install/t_srv_startsrv-aix.html)
- ` Linux ` Starting the server instance (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.install/t_srv_startsrv-linux.html)
- ` Windows ` Starting the server instance (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.install/t_srv_startsrv_win-windows.html)

Operations that were disabled during maintenance mode are reenabled.

Chapter 23. Planning to upgrade the Tivoli Storage Manager server

When a Tivoli Storage Manager fix pack or interim fix becomes available, you can upgrade the server to take advantage of product improvements. Tivoli Storage Manager servers and clients can be upgraded at different times. Ensure that you complete the planning steps before you upgrade the server.

About this task

Follow these guidelines:

- The preferred method is to upgrade the server by using the Tivoli Storage Manager installation wizard. After you start the wizard, in the IBM Installation Manager window, click the **Update** icon; do not click the **Install** or **Modify** icon.
- If upgrades are available for both the server component and the Operations Center component, select the check boxes to upgrade both components.

Procedure

1. Review the list of fix packs and interim fixes. See Technote 1239415 (http://www.ibm.com/support/docview.wss?uid=swg21239415).
2. Review product improvements, which are described in readme files.

   Tip: When you obtain the installation package file from the Tivoli Storage Manager support site (http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Manager), you can also access the readme file.
3. Ensure that the Tivoli Storage Manager version that you upgrade your system to is compatible with other components, such as clients and storage agents. See Technote 1053218 (http://www.ibm.com/support/docview.wss?uid=swg21053218).
4. If your solution includes servers or clients at a level that is earlier than V7.1, review the guidelines to ensure that client backup and archive operations are not disrupted. See Technote 1053218 (http://www.ibm.com/support/docview.wss?uid=swg21053218).
5. Review the upgrade instructions. Ensure that you back up the server database, the device configuration information, and the volume history file.

What to do next

To install a fix pack or interim fix, follow the instructions for your operating system:

- `AIX` Installing a Tivoli Storage Manager server fix pack (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.install/t_srv_inst_fixpack-aix.html)
- `Linux` Installing a Tivoli Storage Manager server fix pack (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.install/t_srv_inst_fixpack-linux.html)
- `Windows` Installing a Tivoli Storage Manager server fix pack (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.install/t_srv_inst_fixpack-windows.html)

Chapter 24. Implementing a disaster recovery plan

Implement a disaster recovery strategy to recover your applications if a disaster occurs and to ensure high server availability.

About this task

Determine your disaster recovery requirements by identifying the business priorities for client node recovery, the systems that you use to recover data, and whether client nodes have connectivity to recovery server. You must also determine how often directory-container storage pools are protected.

For replication, select a hot, warm, or cold recovery site based on your recovery requirements.

Preparing for an outage or system update

Prepare Tivoli Storage Manager to maintain your system in a consistent state during a power outage or system update.

About this task

Ensure that you schedule activities regularly to manage, protect, and maintain the server.

Procedure

1. Cancel processes and sessions that are in progress by completing the following steps:
   a. On the Servers page, select a server for which you want to view processes and sessions, and click **Details**.
   b. Click the **Active Tasks** tab, and select one or more processes, sessions, or a combination of both that you want to cancel.
   c. Click **Cancel**.
2. Stop the server by issuing the **HALT** command:

   ```
   halt
   ```

Completing recovery drills

Schedule disaster recovery drills to prepare for audits that certify the recoverability of the server and to ensure that data can be restored and operations can resume after an outage. A drill also helps you ensure that all data can be restored and operations resumed before a critical situation occurs.

About this task

With a multisite disk solution, use node replication to ensure that data is available on a target server at recovery site and recovery time is fast. When there is an outage, the source server can automatically fail over to a target server for data recovery. If a disaster occurs and the source server is unavailable, client nodes can automatically record information about the target replication server in the client options file. You might need to manually update the client options file for older clients.

Procedure

1. Manually restore data from a target replication server, update the client options file to point to the target replication server. Changes to node replication settings are not required.

2. Configure a client node to store data on a target replication server.

   Restriction: Client nodes that normally back up data to a source replication server cannot back up data to the client nodes that are replicated on the target replication server.

3. Test client data recovery by completing the following steps:

   a. Restore the client system to a similar operating system. Use the same file system names with same amount of file space in the file system

   b. On a system that has enough space for the data, restore the data.

   c. Verify that the client restored successfully. For example, if you restore a virtual machine, verify that the virtual machine powers on and check that the files are available.

Related tasks:

Replicating client node data after a database restore (V7.1.1) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.srv.doc/t_repl_dbrestore.html)

Chapter 20, "Managing replication," on page 129

Chapter 25. Recovering from data loss or system outages

With Tivoli Storage Manager, you can recover data that was lost when a disaster or system outage occurred. You can recover directory-container storage pools, client data, and databases.

Procedure

Use the following recovery methods based on the component you must recover.

Component to recover	Procedure	More information
Directory-container storage pool	To recover directory-container storage pools, complete the following steps: 1. Scan for damaged data extents in the directory-container storage pool by using the **AUDIT CONTAINER** command and specifying the **ACTION=SCANALL** parameter. 2. Repair damaged data extents in the directory-container storage pool by using the **REPAIR STGPOOL** command. **Restriction:** You can only repair a storage pool if the storage pool is protected. 3. Remove damaged data extents by using the **AUDIT CONTAINER** command and specifying the **ACTION=REMOVEDAMAGED** parameter.	"Repairing storage pools" on page 152

Component to recover	Procedure	More information
Client data	**Prerequisites:** • The source replication server, the target replication server, and the client must be at the V7.1 level or later. If any of the servers are at an earlier level, automatic failover is disabled and you must rely on manual failover. Manually configure the client to automatically fail over to the target server for data recovery. If you enabled the client for automated client failover, you can recover the data by using automatic failover function. You can verify that the usereplicationfailover option is either not in the client options file or is set to yes.Recover data from the target server when the source server is unavailable due to an outage by using automatic failover. **Tip:** • • Use the **SET FAILOVERHLADDRESS** command to specify the IP address that a client uses to connect to the secondary replication server during failover, if the address is different from the IP address that is specified for the replication process.	• "Recovering damaged data from a replicated copy" on page 151 • SET FAILOVERHLADDRESS (Set a failover high level address) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_failovrhla_set.html)

Component to recover	Procedure	More information
Database	**Prerequisites:** • To restore the database after a disaster, you must have a copy of the current device configuration file. The device configuration file cannot be recreated. • Ensure that you have a backed up version of the database. Restore the Tivoli Storage Manager database to its most current state or to a specific point in time by using the **DSMSERV RESTORE DB** server utility.	DSMSERV RESTORE DB (Restore the database) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_dsmserv_restore_db.html)

Related reference:

➡ AUDIT CONTAINER (Verify the consistency of database information for a directory-container storage pool) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_container_audit.html)

➡ DSMSERV RESTORE DB (Restore the database) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_dsmserv_restore_db.html)

Restoring the database

You might have to restore the database after a disaster. You can restore the database to its most current state or to a specific point in time. You must have full, incremental, or snapshot database backup volumes to restore the database.

Before you begin

If the database and recovery log directories are lost, re-create them before you issue the **DSMSERV RESTORE DB** server utility. For example, use the following commands: `AIX`   `Linux`

```
mkdir /tsmdb001
mkdir /tsmdb002
mkdir /tsmdb003
mkdir /activelog
mkdir /archlog
mkdir /archfaillog
```

`Windows`

```
mkdir e:\tsm\db001
mkdir f:\tsm\db001
mkdir g:\tsm\db001
mkdir h:\tsm\activelog
mkdir i:\tsm\archlog
mkdir j:\tsm\archfaillog
```

Restrictions:

- To restore the database to its latest version, you must locate the archive log directory. If you are cannot locate the directory, you can restore the database only to a point in time.
- You cannot use Secure Sockets Layer (SSL) for database restore operations.
- You cannot restore a server database if the release level of the database backup is different from the release level of the server that is being restored. For example, if you are using a Version 7.1 Tivoli Storage Manager server and you try to restore a Version 6.3 database, an error occurs.

About this task

Point-in-time restore operations are typically used for situations such as disaster recovery or to remove the effects of errors that can cause inconsistencies in the database. To recover the database to the time when the database was lost, recover the database to its latest version.

Procedure

Use the **DSMSERV RESTORE DB** server utility to restore the database. Choose one of the following methods depending on which version of the database you want to restore:

- Restore a database to its latest version. For example, use the following command:

  ```
  dsmserv restore db
  ```

- Restore a database to a point in time. For example, to restore the database to a backup series that was created on 19 April 2015, use the following command:

  ```
  dsmserv restore db todate=04/19/2015
  ```

What to do next

If you restored the database and directory-container storage pools exist on the server, you must identify inconsistencies between the database and the file system.

1. If you restored the database to a point in time and you did not delay reuse of the directory-container storage pool, you must audit all the containers. To audit all containers, issue the following command:

   ```
   audit container stgpool
   ```

2. If the server cannot identify containers on the system, complete the following steps:

   a. From an administrative client, display a list of containers that the server can identify by issuing the following command:

      ```
      select container_name from containers
      ```

   b. From the file system, display a list of containers by issuing the following command for the storage pool directory on the source server:

      Tip: The storage pool directory is displayed in the command output:

      `AIX`   `Linux`

      ```
      [root@source]$ ls -l
      ```

      `Windows`

      ```
      c:\source_stgpooldir\00>dir
      ```

   c. Compare the containers that are listed on the file system and the server.

    d. Issue the **AUDIT CONTAINER** command and specify the container that is missing from the server output. Specify the **ACTION=REMOVEDAMAGED** parameter to delete the container.

    e. To ensure that the containers are deleted on the file system, review the messages that are displayed.

Related tasks:

↪ Replicating client node data after a database restore (V7.1.1) (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/ com.ibm.itsm.srv.doc/t_repl_dbrestore.html)

Related reference:

↪ AUDIT CONTAINER (Verify the consistency of database information for a directory-container storage pool) (http://www.ibm.com/support/ knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_container_audit.html)

↪ DSMSERV RESTORE DB (Restore the database) (http://www.ibm.com/ support/knowledgecenter/SSGSG7_7.1.3/srv.reference/ r_cmd_dsmserv_restore_db.html)

Recovering damaged data from a replicated copy

If a source replication server is unavailable, you can recover damaged data from a replicated copy that is stored on the target replication server.

Before you begin

The server name that you specify with the **SET REPLSERVER** command must match the name of an existing server definition. It must also be the name of the server to be used as the target replication server. If the server name specified by this command does not match the server name of an existing server definition, the command fails.

Tip:
- Use care when you change or remove a target replication server. If you change a target replication server, client-node data that is replicated is sent to a different target replication server. If you remove a target replication server, client node data is not replicated.

Procedure

1. Verify the replication status of the data on the target server. The replication status indicates whether the most recent backup was replicated to the secondary server.
2. Restore data from a target replication server by setting the source replication server as the target replication server. For example, if you want to set the source replication server as the target replication server, server1, issue the following command:

```
set replserver server1
```

What to do next

When you restore the Tivoli Storage Manager database on a source replication server, replication is automatically disabled. Before you re-enable replication, determine whether copies of data that are on the target replication server are needed.

Related tasks:

➜ Replicating client node data after a database restore (V7.1.1)
(http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/
com.ibm.itsm.srv.doc/t_repl_dbrestore.html)

Repairing storage pools

If a disaster or system outage occurred, you can repair deduplicated data extents in
a directory-container storage pool.

Before you begin

Identify inconsistencies between the database and the directory-container storage
pool by using the **AUDIT CONTAINER** command. By identifying the damaged data
extents in the directory-container storage pool, you can determine what data
extents to repair.

Procedure

To repair a directory-container storage pool, use the **REPAIR STGPOOL** command. For
example, to repair a storage pool, STGPOOL1, issue the following command:

```
repair stgpool stgpool1
```

What to do next

Ensure that no damaged data extents are displayed in the output by using the
QUERY DAMAGED command.

Related reference:

➜ AUDIT CONTAINER (Verify the consistency of database information for a
directory-container storage pool) (http://www.ibm.com/support/
knowledgecenter/SSGSG7_7.1.3/srv.reference/r_cmd_container_audit.html)

➜ QUERY DAMAGED (Query damaged data in a directory-container or
cloud-container storage pool) (http://www.ibm.com/support/knowledgecenter/
SSGSG7_7.1.3/srv.reference/r_cmd_damaged_query.html)

➜ REPAIR STGPOOL (Repair a directory-container storage pool)
(http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/srv.reference/
r_cmd_stgpool_repair.html)

Part 5. Appendixes

Appendix. Accessibility features for the Tivoli Storage Manager product family

Accessibility features help users who have a disability, such as restricted mobility or limited vision to use information technology products successfully.

Accessibility features

The IBM Tivoli Storage Manager family of products includes the following accessibility features:
- Keyboard-only operation using standard operating-system conventions
- Interfaces that support assistive technology such as screen readers

The command-line interfaces of all products in the product family are accessible.

Tivoli Storage Manager Operations Center provides the following additional accessibility features when you use it with a Mozilla Firefox browser on a Microsoft Windows system:
- Screen magnifiers and content zooming
- High contrast mode

The Operations Center and the Tivoli Storage Manager server can be installed in console mode, which is accessible.

The Operations Center help system is enabled for accessibility. For more information, click the question mark icon on the help system menu bar.

Vendor software

The Tivoli Storage Manager product family includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for the accessibility information about its products.

IBM and accessibility

See the IBM Human Ability and Accessibility Center (http://www.ibm.com/able) for information about the commitment that IBM has to accessibility.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
United States of America

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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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 want 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
US

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.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

```
Portions of this code are derived from IBM® Corp. Sample Programs.
© Copyright IBM® Corp. _enter the year or years_. All rights reserved.
```

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Linear Tape-Open, LTO, and Ultrium are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability
> These terms and conditions are in addition to any terms of use for the IBM website.

Personal use
> You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use
> You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights Except as expressly granted in this permission, no other permissions,

licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

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, specific 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 "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

Glossary

A glossary is available with terms and definitions for the IBM Tivoli Storage Manager family of products.

See Tivoli Storage Manager glossary (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/tsm/glossary.html).

To view glossaries for other IBM products, see http://www.ibm.com/software/globalization/terminology/.

Index

IBM ®